

(19) World Intellectual Property Organization  
International Bureau(43) International Publication Date  
21 June 2001 (21.06.2001)

PCT

(10) International Publication Number  
WO 01/44935 A1

(51) International Patent Classification<sup>7</sup>: G06F 9/445 (74) Agent: PLAÇAIS, Jean-Yves; Cabinet Netter, 40, rue Vignon, F-75009 Paris (FR).

(21) International Application Number: PCT/IB99/02006 (81) Designated States (national): AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(22) International Filing Date: 15 December 1999 (15.12.1999) (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

(25) Filing Language: English (26) Publication Language: English

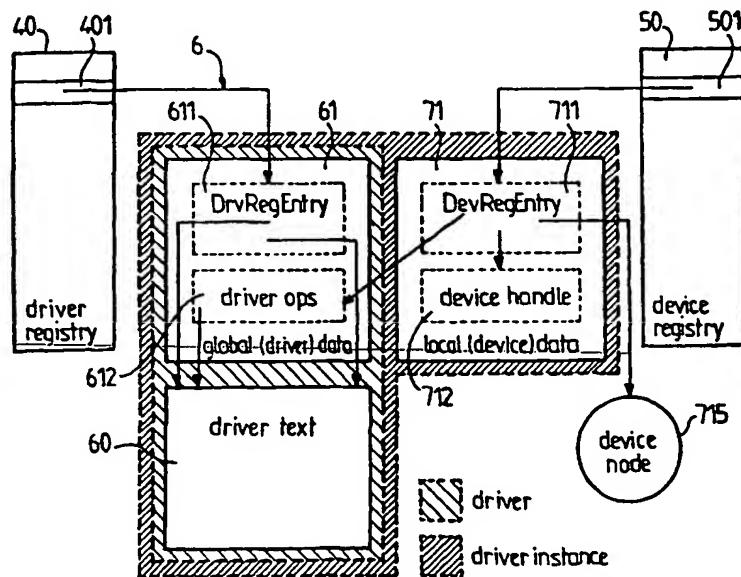
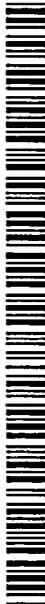
(71) Applicant (for all designated States except US): SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, Palo Alto, CA 94303 (US).

(72) Inventors; and (73) Inventors/Applicants (for US only): RUGET, Frédéric [FR/FR]; 30, avenue Duquesne, F-75007 Paris (FR). GROUZDEV, Vladimir [RU/FR]; 6, rue Alexandre Bertrand, F-78100 Saint Germain en Laye (FR). LE-SCOUET, Eric [FR/FR]; 144, boulevard Voltaire, F-75011 Paris (FR).

Published:  
— With international search report.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: COMPUTER SYSTEM WITH AN IMPROVED DEVICE AND DRIVER FRAMEWORK



WO 01/44935 A1

(57) Abstract: Each driver has a global driver section (6) comprising driver program code (60) and global driver data (61), both defining driver operations associated with a corresponding class of devices. Each device has an associated driver instantiation (71) comprising local driver data, having an internal pointer to the global driver section of the instantiated driver, and possibly at least one external pointer to another driver instantiation. A device registry (50) forms a list of device registry entries (711), each corresponding to a respective one of the driver instantiations. Driver client code may perform operations on a given device by fetching in the device registry (50) a pointer to the driver instantiation (71) associated with the given device, and then use operations provided by the associated driver instantiation, directly or using further operations reached in other driver instantiations by using iteratively the external pointer.

COMPUTER SYSTEM WITH AN IMPROVED DEVICE AND DRIVER FRAMEWORK

5

This invention relates to computer technology.

Real time embedded operating systems are of increasing interest in various applications. An example is the operating system (OS) named ChorusOS (a product and trademark of SUN MICROSYSTEMS), which is flexible as to the hosting hardware. This means that it may have to be installed and configured on a variety of computer platforms.

10 15 For being installed on a particular machine ("target"), an OS like ChorusOS has to be prepared in accordance with the target characteristics, including its main board, the corresponding board support package (BSP), and its specific drivers. Usually, the target will include additional devices, 20 also with their own drivers.

Thus, in view of the desired flexibility, a number of different drivers may have to be loaded, depending upon the systems. The ability to hot swapping a device is also 25 desirable. Hot swap means maintaining the system functions for the time interval extending from unplugging an existing device to plugging in a new device in substitution of the unplugged existing device. In other words, hot swap includes both hot unplug and hot plug, and continue to play while 30 maintaining system functions in each case.

This invention intends to improve the driver and device framework in such a system.

35 Another object of this invention is to provide a very scalable way to configure/initialize a set of device drivers to operate a hardware platform.

Another object of this invention is to provide a driver and device framework adapted to range from low memory footprint static software configurations, where a built-in driver is statically bound to each supported device, to very large 5 hardware and software configurations where the set of physical devices is not known a priori.

A further object of this invention is to enable devices to be dynamically added in/removed from the system, and/or to 10 enable devices to be dynamically probed.

Still another object of this invention is to enable the downloading of drivers from the network.

15 The invention applies to a computer system comprising a processor, a memory, devices, and a program memory area, for storing an operating system, drivers, and driver client code (applications).

20 Each driver has a global driver section comprising driver program code and global driver data, both defining driver operations associated with a corresponding class of devices.

For each operative one of the devices, there is provided an 25 associated driver instance comprising local driver data. The local data may have an internal pointer to the instantiated driver. They may also have at least one external pointer to another driver instance. The local driver data may still further comprise local device data, defining a current 30 condition of the associated device.

A device registry forms a list of device registry entries, each corresponding to a respective one of the driver instances. A driver client code is thus capable of performing 35 operations on a given device by:

- . fetching in the device registry a pointer to the driver instance associated with the given device,
- . performing the given device operations provided by the associated driver instance, directly or using further

operations reached in other driver instances by using iteratively said external pointer.

5 The invention may also be defined as a method of managing drivers, or a driver framework, in a computer, comprising the steps of:

10 a. providing each driver with a global driver section comprising driver program code and global driver data, both defining driver operations associated with a corresponding class of devices,

15 b. for each operative one of the devices, providing an associated driver instantiation comprising local driver data, having an internal pointer to the instantiated driver,

c. providing a device registry, forming a list of device registry entries, each corresponding to a respective one of the driver instantiations.

20 Thus, a device call may be directed to the device registry, to reach the driver instantiation corresponding to the device being called.

25 Preferably, provision is made for device tree data, defining a tree representation of at least some of the devices. In the tree, each device node represents a device and its device characteristics, and identifies at most one driver capable of managing that device. The device registry entry may further comprise a node identifier, defining an associated node in the device tree.

30 On another hand, the global driver section may comprise code adapted to implement at least one of the following special operations, non limitatively: self-registering the driver in the driver registry, identifying devices the driver can serve, creating a node in the device tree, binding a device with a driver, creating a new driver instance, unloading a driver.

According to another feature, there is provided for a driver registry, forming a list of driver registry entries, each comprising at least one pointer to the special operations.

The driver registry may directly or indirectly define device-  
5 independent driver characteristics, comprising a driver identifier and a parent driver class identifier. In particular drivers, the global driver section may comprise code adapted to explore the driver registry for finding drivers matching the class of the particular driver.

10

The invention also includes the software code portions being used, including individual driver code, and portions of code in the operating system, or in its nucleus, which implement the functionalities of this invention.

15

Other features and advantages of the invention will appear in the detailed description below and in the appended drawings, in which:

- Fig. 1 schematically illustrates a hardware computer structure in which the invention may be practiced;
- Fig. 2 schematically illustrates a device driver system for the hardware computer structure of Fig. 1;
- Fig. 3 illustrates a device driver framework in accordance with an embodiment of the invention for the hardware computer structure of Fig. 1;
- Fig. 4 illustrates a device registry and a driver registry, cooperating with a global driver and an instantiation of that global driver;
- Fig. 5 illustrates an exemplary device tree;
- Fig. 6 illustrates a corresponding exemplary driver tree;
- Fig. 7 is a flowchart illustrating a driver registration;
- Fig. 8 is a flowchart illustrating the main steps of bus driver initialization;
- Fig. 9 is a flowchart illustrating step 910 of Fig. 8 in more detail;
- Fig. 10 is a flowchart illustrating step 920 of Fig. 8 in more detail;
- Fig. 11 is a flowchart illustrating step 930 of Fig. 8 in more detail;

- Fig. 12 is a flowchart illustrating the steps of leaf driver initialization;
- Fig. 13 is a flowchart illustrating the steps of processing a load handler;
- 5 - Fig. 14 shows table T1, useful for the detailed description;
- Fig. 15 shows table T2, useful for the detailed description;
- Fig. 16 shows table T3, useful for the detailed 10 description;
- Fig. 17 shows table T4, useful for the detailed description;
- Fig. 18 shows table T5, useful for the detailed description; and
- 15 - Fig. 19 shows table T6, useful for the detailed description.

Additionally, the detailed description is supplemented with Exhibits, in which:

- 20 - Exhibit A contains pseudocode of various routines in C type language; and
- Exhibit B contains comments on the routines of Exhibit A, using a corresponding paragraph numbering.

25 The Figures and Exhibits include graphical and/or code information, which may be useful to define the scope of this invention.

30 Making reference to software entities imposes certain conventions in notation. For example, in the detailed description, the quote sign " may be used as a string delimiter wherever deemed necessary for clarity (e.g. "actor"). Square brackets may be used to frame the optional portion of an expression (e.g. "[OS] archive").

35 The detailed description hereinafter refers to platforms based on the ChorusOS operating system, and uses the corresponding terminology. This is exemplary only and it should be understood that the invention is applicable to a

variety of computer platforms. Generally, ChorusOS has the following abilities:

- work with various processors and various main boards, adapted to specific products;
- 5 - be loadable in RAM from a so called "[OS] archive" or "image file", stored e.g. in persistent memory or as a file on the hard disk, depending upon the actual platform architecture;
- define independent processes (named "actors"), exchanging
- 10 messages.

The exemplary hardware computer structure of Fig. 1 comprises a CPU 110, having a CPU Bus 112, and RAM 115 connected thereto. A PCI Bus Bridge 118 is interposed between CPU Bus 112 and PCI Bus 122. PCI Bus 122 in turn is connected to an Ethernet device 126, an SCSI Host Bus Adapter 124, and an ISA Bus Bridge 128. SCSI Host Bus Adapter 124 is connected to SCSI Bus 130, which is connected to one or more hard disks 132, and one or more CD-Rom drives 134. ISA Bus 135 is connected to one or more UART devices 136, and to one or more floppy drives 138.

As known, the devices necessitate pieces of code named drivers to be able to correctly operate. In accordance with 25 one aspect of this invention, the tendency is to systematically have one driver for each device, with the driver services being those required for the device.

Accordingly, the device driver framework has the structure 30 shown in Fig. 2. The micro kernel basic services, including CPU Interrupt management and memory management, form a module 205, with PCI Bus Bridge driver 218 connected thereto. Module 218 in turn receives connection from SCSI Host Bus Adapter driver 224, and from ISA Bus Bridge driver 228, as well as 35 from an Ethernet peripheral driver 226. SCSI Host Bus Adapter driver 224 has one or more hard disk drivers 232, one or more CD-Rom drivers 234. ISA Bus Bridge driver has one or more UART drivers 236 and one or more floppy drivers 238.

The Interface layering will now be described with reference to Fig. 3. The interfaces specified by ChorusOS for the device drivers implementation are split logically in two sets:

- 5 - Drivers/Kernel Interface (DKI), which comprise a set of services provided by the microkernel 205 for driver implementation. In other words, all DKI services are implemented in the microkernel, and used by device drivers. Typically drivers for buses and devices connected to the CPU local bus will use all DKI services.
- 10 - Device Drivers Interface (DDI), which comprises a set of services provided by drivers for the driver clients implementation. In fact, a driver client may itself be a driver. For instance a device driver is typically a bus driver client. Thus, all DDI services are implemented in a driver, and used by the upper layer drivers, or directly by drivers client applications.
- 15
- 20 The DKI interface defines all microkernel services provided for drivers implementation. These services are split in two categories. Common DKI services 207 have the same API for all platforms and processors, and are usually used by all kind of drivers, whatever the layer level is. Processor family specific DKI services 206 may have a different API for each different processor family and are usually used only by the lowest-level drivers, i.e. drivers for buses and devices which are directly connected to the CPU local bus 112 (Fig. 1).
- 25
- 30 The Processor family specific DKI services 206 cover: Processor interrupts management, Processor caches management, Processor specific I/O services and Physical to virtual memory mapping. Some or all of these services exist, and may be implemented in known fashion. These DKI services operate with the CPU bus (the root of the device tree) to some extent like DDIs with devices; there is a DKI service for each particular CPU bus function, and the design of the DKI services is rendered modular, thus improving the portability.
- 35

Fig. 3 also shows that the other drivers have interfaces arranged as DDI services as follows:

- DDI 219 for PCI bus bridge driver 218;
- DDI 225 for SCSI host bus adapter driver 224;
- 5 - DDI 229 for ISA bus bridge driver 228; and
- DDI 240 diagrammatically showing the DDIs for all other end drivers, communicating with driver clients 250 in application programs.

10 Fig. 4 shows a driver registry 40, having a particular driver record 401, and a Device Registry 50, having a particular device record 501.

Many drivers exist. Only one is shown in fig. 4 for clarity.  
15 A driver 6 has driver program code ("driver text") 60 and global driver data 61, comprising:

- a driver operations section ("driver ops") 612, adapted to provide services or functions, using driver program code 60, and
- 20 - a driver registry entry section ("DrvRegEntry") 611. Driver record 401 has a pointer to "DrvRegEntry" 611 of the particular driver 6 being considered. For initialization and shut-down, the "DrvRegEntry" 611 may optionally define entry points to device independent routines in the driver, e.g. 25 "drv\_probe", "drv\_bind", "drv\_init" and "drv\_unload", to be discussed hereinafter.

As shown in Fig. 5, the exemplary device structure of Fig. 1 can be rearranged into a tree, having a root 500 constituted 30 by the CPU Bus. The CPU and RAM appear as leaf nodes 510 and 511. PCI Host Bus Bridge is an internal node or "bus/nexus" 512, which has subnodes SCSI HBA 521, and ISA Bus Bridge 522, and leaf node 520 for an Ethernet device. In turn, SCSI HBA 35 521 node has subnodes 531 for the hard disk and 532 for a CD-ROM drive. Also ISA Bus Bridge node 522 has subnodes 534 for the UART and 536 for a floppy disk drive. The peripheral devices are leaf nodes (more briefly "leaves") in the tree, here CPU, RAM, Ethernet, Hard disk, CD-Rom, UART and floppy.

As above indicated, e.g. one or more hard disks may be provided. In the case of two hard disks, there are of course two hard disk leaf nodes connected e.g. to SCSI HBA nexus 521.

5

In this invention, driver 6 of fig. 4 is a generic driver object or "global" driver section, which is not used as such by devices. In accordance with the concepts of object programming, the devices needing the services of the 10 particular driver 6 will use "instances" of the object (object class) defined by driver 6.

Such an instance basically comprises "local" driver data 71, defining the instance, together with the global driver 15 section. The local driver data 71 comprises data forming an "internal" pointer to the driver object 6. "Internal" means that the pointer connects the driver instance to the global driver section within the driver. In other words, a driver instance comprises its local driver data plus the global 20 driver section. Preferably, the local driver data 71 also comprises data forming an "external" pointer to at least one other driver, typically the parent driver (the root's parent driver are the DKI services).

25 The local driver data 71 may also include local device-oriented data, defining, in section 712:  
- data defining the current condition of the local device, and  
- a device handle, for the local device associated with the 30 driver instance.

Consideration is now given to the case of a driver instantiation for a leaf device in Fig. 4, i.e. one of the devices directly communicating with driver clients in Fig. 3. 35 Such a leaf driver instance is also associated with a "Device registry entry". In Fig. 4, the device registry entry is shown as a section 711 in "local" device data 71 of the driver instance (However, the Device registry entry may alternatively be associated with the corresponding record 501

in device registry 50). In any case, each item in the Device Registry 50 in fig. 4 has a corresponding "DevRegEntry" structure 711, which is exemplified in pseudocode in Exhibit A-1, and commented in Exhibit B-1.

5

Device registry 50 forms the basis of a client-to-[leaf]-driver linking mechanism. For subsequently reaching a bus/nexus driver, a child-to-parent-driver linking mechanism is used.

10

As shown in Fig. 6, the device related software may also be defined and split in accordance with a tree structure, in parallelism with the device tree. The tree includes a root node 550, which has a CPU management leaf 560, for example for interrupts, and a memory management leaf 561, plus the PCI HBB driver nexus 562. This nexus in turn has an Ethernet driver leaf 570, and two sub-nexus:

- SCSI HBA driver 571, with hard disk driver leaf 581 and CD-ROM driver leaf 582, and
- ISA BB driver 572, with UART driver leaf 584 and floppy driver leaf 586.

As reflected by the dashed line frame in Fig. 6, a portion of the device related software, e.g. the CPU management and memory management, is part of the nucleus of the OS, rather than of the driver framework as such. When restricting the driver tree of Fig. 6 to the drivers per se, its root is the PCI HBB driver nexus 562.

30 The device tree of Fig. 5 is preferably represented as such in the code. By contrast, the driver tree of Fig. 6 is virtual, i.e. is defined only by the fact each device in the device tree of Fig. 5 should normally have a corresponding driver instantiation.

35

A computer in operation will now be considered. It is assumed that driver registry 40 contains drivers as required to satisfy the following conditions:

- each bus/nexus node in operation has found at least one driver matching its own needs, and a driver instance for that nexus node has been created;
- each leaf node in operation has been provided with at least 5 one driver matching its own needs, and a driver instance for that leaf node has been created;
- each device having a leaf driver instance has a Device registry entry, accessible through device registry 50.

10 When a driver client 250 wishes to access a particular leaf device, e.g. hard disk 132, to perform a given operation on it, the process is as follows:

- the device is searched for in device registry 50, finding e.g. record 501, and the DevRegEntry 711 in the driver 15 instance 232 corresponding to that device;
- in turn, DevRegEntry 711 points to driver operations in 612, which are thus made available to the device;
- the adequate device handle is given by the data in 712; this handle is a pointer to associate internal data of the 20 driver instance with the device using these data;
- a pointer to the corresponding device 715 is also available (or, if the device tree structure is used at this stage, a pointer to the corresponding device node in the device tree);
- if the operation needs, another driver is called to help in 25 executing that operation, using further external pointer(s); this is repeated as necessary.

Thus, in a computer in operation, all that is necessary to 30 process a call from an application (client 250) to a particular leaf device may be made available by the cooperation of device registry 50 with the corresponding local data 71.

A more detailed description will now be given, including the 35 driver initialization (and, optionally, shutdown).

The common DKI services 207 include a "DKI thread", launched by the microkernel at initialization time. The DKI thread is part of the nucleus (or nucleus actor in ChorusOS) of the

operating system. The DKI thread serves as a mechanism to synchronize same calls to DKI and DDI services issued from different threads at drivers initialization time, and also at runtime (if required, e.g. for initialization and shutdown of 5 drivers). Thus, using the DKI thread avoids the need for other synchronization mechanism (locks) in the driver implementations themselves.

Normally, all calls related to initialization / shutdown of 10 the drivers are done implicitly in the context of the DKI thread. Since the corresponding driver routines are called directly from the DKI thread, there is no need for code taking care of synchronization. However, a driver may explicitly use DKI thread services to ensure synchronization 15 in certain special cases, like hot-pluggable device drivers, or deferred initialization and shutdown process, which are described later on.

The basic calls to the DKI thread are exemplified in 20 pseudocode in section A-2 of Exhibit A, and described in more detail in section B-2 of Exhibit B. The contexts in which a caller is allowed to invoke each service are shown in the table T1 of Fig. 14. The possible calling contexts in the tables are:

- 25 - Base level: the service may be called from any thread base level context.
- DKI thread: the service may be called from DKI thread context.
- Interrupt: the service may be called from interrupt 30 context, i.e. from the context of an interrupt handler.
- Blocking: the service may block and preempt the caller.

Other features of the Common DKI 207 are: Device tree, Driver 35 registry, and Device registry. Common DKI 207 may also include other services not related to this invention: General purpose memory allocation, Timeout, Precise busy wait, Special purpose physical memory allocation, System event management, Global interrupts masking, and Specific I/O services. Again, there is a DKI service for each particular

common function, i.e. the design of the DKI services is rendered modular.

5 The device tree is a data structure which provides a description of the hardware topology and device properties. Tree services are available from a device tree API, exemplified in pseudocode in section A-3 of Exhibit A, and described in more detail in section B-3 of Exhibit B. In these examples, "DevNode" is an abstract type designating a 10 device node object, and is opaque for the driver. More precisely:

- functions A-3.1 described in B-3.1 define a low-level API section related to the device tree browsing;
- functions A-3.2 described in B-3.2 define a low-level API 15 section related to the device tree topology modification;
- functions A-3.3 described in B-3.3 define a low-level API section related to the properties of the device tree nodes;

20 From the above low-level functions, an API section offering Device tree high-level services is built, implementing some useful services for building and searching the device tree, to avoid having low level code implemented in all device tree users. Functions A-3.4 described in B-3.4 define these Device 25 tree high-level services. Table T2 of Fig. 15 specifies the contexts in which a caller may be allowed to invoke each low-level or high level service. Building the above functions on a tree structure is considered accessible to men skilled in the art, using Exhibits A and B as necessary.

30 The hardware topology of the devices is specified in terms of parent/child relationships. Each device node in the tree is associated with device properties, which are device specific. A device property is a "name/value" pair. The property "name" is e.g. a null terminated ASCII string. The property "value" 35 is a sequence of bytes specified by the "length/address" pair. The property value format is property specific and may have to be standardized between the given producer and its consumers. For instance, certain device node properties are related to the bus resources required and/or allocated for

the device (e.g. interrupt lines, I/O registers, DMA channels). Such properties are preferably standardized, in order to be understood by the bus driver as well as by device drivers connected to the given bus.

5

Preferably, a device registry microkernel module ("device registry API") implements a data base of driver instances servicing devices currently supported by the operating system. The device registry data base/has been populated by 10 drivers at device initialization time. The device registry data base is accessed by driver clients in order to obtain a pointer to the driver instance servicing a given (logical) device.

15 Device registry management services or functions (the API) are exemplified in pseudocode from A-4.1 to A-4.8 and described in corresponding sections B-4.1 to B-4.8. Table T4 in Fig. 17 specifies the contexts (as defined above) in which a caller is allowed to invoke each service. Note that only 20 the "svDeviceLookup", "svDeviceRelease" and "svDeviceEntry" microkernel calls should be used by driver clients 250. The rest of the API is dedicated to device drivers.

25 Preferably, the driver code in each driver 6 contains a "main" routine, for performing self-registration of the driver in Driver registry 40.

Driver registry module 40 implements a data base of drivers, which is useful for driver initialization and shutdown. Each 30 item in the Driver Registry 40 in fig. 4 has a pointer to the "DrvRegEntry" structure 611, which is exemplified in pseudocode in A-5, and commented in B-5.

35 Driver registry management services or functions (an API) are exemplified in pseudocode from A-6.0 to A-6.6 and described in corresponding sections B-6.0 to B-6.6. Table T3 in Fig. 16 specifies the contexts (as defined above) in which a caller is allowed to invoke each service.

Each driver's "DrvRegEntry" structure 611 may include one or more driver management oriented functions, or driver entry points. The three following functions may be invoked by a bus/nexus driver when the "bus\_class" specified in the driver 5 registry entry matches the bus/nexus class:

- "drv\_probe", which is able to detect device(s) residing on the bus, with a view to create a device node for each such device.
- "drv\_bind", which checks whether a detected device may be 10 serviced by the driver.
- "drv\_init", which creates an instance of driver servicing the "bound" device, and performs other steps as necessary for full driver/device cooperation.

15 A fourth function, "drv\_unload", may be invoked by the driver registry module when an application wishes to unload a driver component from the system.

In "DrvRegEntry" structure 611, the "bus\_class" and 20 "bus\_version" fields specify a parent bus/nexus API required for the driver. Such an API is provided to the driver as a structure of indirect functions implementing the bus API service routines. For instance, the "pci" bus API is specified by the "PciBusOps" structure, the "isa" bus API is 25 specified by the "IsaBusOps". When a bus driver invokes the "drv\_probe" or "drv\_init" driver's routine, it gives a pointer to the structure of the bus service routines, in its driver global section. The structure type corresponds to the "bus\_class" field value.

30 Accordingly, the following will refer to a "bus\_ops"+"bus\_id" pair, which specifies the parent driver of a current driver. "bus\_ops" points to a structure of service routines implementing a bus driver API in the parent driver. Such a 35 structure is bus class specific, i.e. several such structures may exist within the "driver\_ops" of the parent driver, implementing several bus classes. The structure being aimed at by "bus\_ops" corresponds to the "bus\_class" specified by the driver registry entry of the current driver. "bus\_id" is

an opaque for the current driver. It must be passed back to the parent bus driver, e.g. when the "open" bus service routine is invoked.

5 This invention uses mechanisms implemented in order to support the dynamic loading/unloading of a driver component, the dynamic discovering/enumeration of devices residing on a bus, the dynamic initialization/shut-down of a driver instance. These issues concern bus/nexus as well as device  
10 drivers.

The way the computer starts operation will now be considered.

15 It should be reminded that the device tree structure is accessed in the DKI thread context only, to ensure access synchronization. In many cases, driver routines which typically examine/modify the device tree are already invoked by the DKI thread. For instance, the driver probing, initialization and shut-down code is executed by the DKI  
20 thread. In a case when a driver needs to access the device tree at run time, a driver routine which uses the device tree API must be explicitly called in the DKI thread context. The same rule is applied to the driver clients: a driver client must always switch to the DKI thread context when needing to  
25 access the device tree structure.

The device tree data structure may be built either statically or dynamically:

- in the static case, the device tree is populated by the OS  
30 booter. For instance, the OS booter may include a pre-defined sequence of device tree function calls. Another possibility for the OS booter is to build the device tree from a hardware description provided by firmware.
- in the dynamic case, the device tree is populated at system  
35 initialization time using an enumeration/probing mechanism. The device tree is populated by propagating from parent to children.

It is possible to combine both methods: for example, an initial (non complete) device tree may be provided by the OS booter, and will later be completed dynamically using an enumeration/probing mechanism. The device tree structure may 5 also be modified (extended/truncated) dynamically at runtime, because of the hot-plug insertion and /or removal (e.g. PCMCIA cards).

10 The following is a typical computer starting sequence (as far as devices and drivers are concerned).

15 The first phase is the driver registration time (800, Fig. 7), initiated by the OS nucleus. At step 802, a driver component is downloaded in some way: it may be built-in to 20 the OS bootable image, or it may be downloaded dynamically as a supervisor actor using e.g. the "afexec" system call of ChorusOS. At step 804, the system calls the "main" routine which performs the self-registration: it invokes the "svDriverRegister" microkernel call, passing as an argument 25 a "DrvRegEntry" data structure which specifies the driver component properties. The driver is now added in driver registry 40 at step 806. Its DrvRegEntry structure is now accessible to the OS nucleus, e.g. through a pointer. Steps 804 and 808 continue until all downloaded drivers appear in 30 driver registry 40 (step 810).

35 After the driver components have self-registered, the future driver management is mostly driven by the parent bus/nexus driver of each driver component, using the properties specified in the "DrvRegEntry" structure of that driver.

The second phase may be called "enumeration/probing". At that time, the OS nucleus starts execution of an iterative and/or recursive process, using the device tree, the device 35 registry, the driver registry and certain of the driver operations. In the exemplary embodiment, this process forms a cascade of internested processes of "bus/nexus driver initialization".

Generally, an elementary phase of the process is as follows:

s0. a bus-nexus driver, bound to its device, is considered,

s1. the bus-nexus driver is initialized as a driver (like a leaf driver is initialized),

5 s2. in the case of a bus-nexus driver, the driver initialization additionally comprises (normally):

s21. a "probing" step for finding child devices of the current bus-nexus driver,

s22. a "binding" step for matching such child devices with 10 corresponding drivers, and

s23. an "init" step for initializing the driver/device pairs, by creating driver instances as necessary.

The iteration is due to the fact that, whenever step s23 15 initializes a bus nexus driver, a nested similar process is initiated, starting again at step s0, as it will be seen.

It should be reminded that, in Fig. 6, the root node 550 is the DKI services, which is part of the nucleus. Initiation of 20 the second phase begins with a "bus/nexus driver initialization", applied to the (driver) root, whose parent is the DKI thread. In the example, the driver root is the PCI Bus Bridge driver 218. The corresponding node in the device tree is the PCI bus node. The required condition and resources for 25 the PCI bus to operate are supposed to be prepared by the nucleus.

The "bus/nexus driver initialization" will now be described in general, i.e. whatever the bus nexus is. Fig. 8 shows the 30 main steps of such a Bus Driver Initialization Routine 900 (or "BDIR").

The first step 902 of the initialization of a given bus/nexus driver is similar to the leaf device driver initialization 35 process of Fig. 12, and will be described in connection therewith. At this point, it is assumed that, after step 902, the given bus/nexus driver is able to access its hardware (e.g. internal bus bridge registers), that it has established connection to its parent bus/nexus driver, that it can use

services implemented by its parent driver. (Initially, for the root driver, here the PCI bus, this has been done by the nucleus).

5 Then comes step 910 of "probing", which may result into the creation of child device nodes in the device tree. This will be described in more detail with reference to Fig. 9.

10 Step 919 deals with allocating resources to the newly entered devices at step 910, as defined in step 918 of Fig. 9 (to be described). Each device node specifies bus resource properties required for the associated device: some of the required bus resources may be hardwired (e.g. fixed IRQ); other required bus resources may be configurable, and subject  
15 to some constraints (e.g. device address decoder constraints). At step 919, the bus/nexus driver iterates through existing child nodes in order to allocate configurable resources (taking the constraints into account) and to check possible resource conflicts. If a resource  
20 conflict is detected, the bus/nexus driver behaviour is implementation specific. The bus/nexus driver should never activate any driver on a node for which bus resources are not allocated successfully.

25 Then, step 920 performs a driver-to-device binding, which will be described with reference to Fig. 10. Finally, at step 930, driver instances are created in the child device nodes. This will be described with reference to Fig. 11.

30 Fig. 9 illustrates the probing step 910 in more detail. A "current" bus/nexus driver (911) searches in the driver registry data base for all drivers satisfying the following condition: the parent device class specified in the driver being searched must match the class of "current" bus/nexus,  
35 or one of its classes (913). The drivers being found (914) will form the first level of children of the "current" bus/nexus driver (the one considered in 911). For each such driver (implementing the "drv\_probe" entry point), step 918 immediately executes another "drv\_probe".

The pseudocode given in Exhibit A.7 and commented in Exhibit B.7 illustrates a typical probing loop in a bus/nexus driver, creating device nodes as required. "drv\_probe" is called with three arguments: "dev\_node", specifying the parent device 5 node and a "bus\_ops" and "bus\_id" pair specifying the parent device driver. The "drv\_probe" routine is optional. If it is absent, the "drv\_probe" field in "DrvRegEntry" is set to NULL.

10 In general, the operation of a probe routine may be summarized as follows:

- The probe routine should avoid creating redundant nodes.
- The probe routine specifies the following device node properties,

15 . a physical device ID, enabling the bus driver to find the appropriate device driver for this device node. Preferably, the device ID is bus class specific. For instance, on PCI bus, the device ID is the vendor/device IDs pair.

20 . resource requirements, so that the bus driver can reserve resources required to initialize the device.

The probe routines will now be considered in more detail. In fact, there are two kinds of probe routines:

- generic (i.e. bus class specific only), like a self-25 identifying bus (e.g. PCI) enumerator;
- device specific (i.e. bus class and device specific), like a device probing code on an ISA bus.

Accordingly, child driver offering multiple probe routines 30 for a given parent bus may be found in the driver registry. They may be processed in any desired order.

In addition, the probe routines may be invoked at run time when, for example, a device insertion is detected on the bus. 35 In the latter case, the probe routines must be extremely careful about active device nodes, i.e. existing device nodes for which the device drivers have been already started and may be already in use.

The following rules are preferably respected by generic and device specific probe routines:

- The generic and specific probe routines may access the device hardware only through the bus service routines. The bus resources needed to access the device hardware (e.g. I/O registers) may be allocated through the bus service routines ("resource alloc"). This prevents the probe routine to access hardware which is currently in use. Upon unsuccessful probing, the used hardware resources is released through the bus service routines ("resource free").
- 5
- 10
- 15
- 20
- 25
- No probe routine (whether generic or device specific) may delete an active device node, or modify its properties (An active device node is a node for which a device driver is already running; such nodes are flagged with the "active" property).
- Device specific probe routines have higher priority than generic ones, i.e. device specific probe routines are allowed to override properties in an existing node or to delete an existing node, when inactive, while generic probe routines are not.
- No probe routine may create redundant nodes. When a probe routine is to create a device node, then it verifies that no other node exist for this device. Probing is normally avoided if the probing routine may create redundant nodes.

Implementing such probing mechanisms within a driver is considered accessible to men skilled in the art.

A preferred implementation of the device-to-driver binding step 920 is shown in more detail on Fig. 10.

A "current" bus/nexus driver (921) searches in the driver registry data base for all drivers satisfying the following condition (923): the parent device class specified in the driver being searched must match the class of "current" bus/nexus, or one of its classes (and implement the "drv\_bind" entry point). Each of the drivers being found (924) may now be bound to known device child nodes, found by the probing, and appearing in the device tree. Steps 926 and

927 scan through these child device nodes, and, in each, the "drv\_bind" routine is invoked at step 928. When all children for all matching drivers have been seen, step 924 goes to the end.

5

The pseudocode given in Exhibit A.8 and commented in Exhibit B.8 illustrates a typical binding loop in a bus/nexus driver.

10 The driver-to-device binding mechanism may have various implementations:

15 - a very simple bind routine may be implemented by the child device driver itself. Such an implementation is device specific, and applies only to devices known by the driver as compatible with a reference device for which the driver is written.

20 - on the other hand, the system may need to support after-market hot-plug devices, and consult some network look-up service to locate the driver for a new device. In such a case, it may be preferable to use a generic separate binder driver, dedicated to implementing a smart driver-to-device mapping and a driver component downloading (without the probe and init routines). Of course, such a (generic) binder will be treated like the other driver components.

25 A preferred implementation of step 930 is shown in more detail on Fig. 11.

30 At step 932, each bus/nexus driver explores its child device nodes in the device tree. It tests whether any of these devices is bound to a driver, but not active (933). This means that the driver has not yet been initialized, and should be started and instantiated for that device. Then, in the described embodiment, a complete checking is again performed: for any device verifying the condition, the driver 35 registry is again explored (935); if a driver (or a further one) is found at step 936, then step 937 checks whether a) the required bus class of that found driver matches one class of the bus/nexus, and b) the driver name matches the one bound to the device node being considered. If so, step 940

calls the "drv\_init" routine of that driver. The "drv\_init" is called with three arguments. The "dev\_node" argument specifies the device node for which a device driver instance should be created. If the drv\_init succeeds in activating the 5 device, test 939 returns to 933. Otherwise, or if step 937 indicated misfit, control is returned to step 936. If no adequate driver is found, test 936 returns control to step 933, for processing another child device, if any.

10 The above described complete checking is of interest because the same routine may be used for "hot" device and driver management as well. When used at initialization only, the routine may be simplified.

15 Preferably, when a driver instance is activated by a bus/nexus driver, the child driver establishes a connection to its parent bus/nexus driver, while specifying a call-back event handler. The parent bus/nexus driver may use the call-back event handler mechanism e.g. to shut down the child 20 driver instance later, if required.

Step 940 will now be described in further detail.

If the child is a leaf child node, the process to be 25 described with reference to Fig. 12 is executed, and control is given to step 939 of Fig. 11.

If, by contrast, the child is a bus-nexus, the process of Fig. 8 is resumed with the child. This means that a new BDIR 30 (bus driver initialization routine) is recursively called, to be nested within the current one. The status of the current BDIR is stored before initiating the new one. This new BDIR may in turn result into a further BDIR, and so on, until reaching a "leaf-only" BDIR, where the child level includes 35 leaf nodes only. When such a leaf-only BDIR terminates, by the "N" output of step 933, control is returned to the immediately preceding BDIR, until it terminates in turn. This is repeated until the initial BDIR terminates, indicating the whole device tree has been scanned.

The pseudocode given in Exhibit A.9 and described in Exhibit B.9 illustrates a typical initialization loop in a bus/nexus driver.

5     The purpose of "drv\_init" is to create an instance of driver servicing the given device, to perform a device hardware initialization and to register the driver instance in the device registry (in the example, for leaf only). Typically, "drv\_init" would read the device and bus node properties in  
10    order to obtain the assigned bus resources and tunable parameters related to the bus/device.

The "drv\_init" routine is optional. In its absence, the "drv\_init" field is set to NULL. A typical example of a probe-only driver is a self-identifying bus enumerator (e.g. PCI enumerator) which is implemented as a driver. Such a driver has the "drv\_probe" routine which enumerates devices residing on the bus and creates device nodes, but needs no "drv\_init" routine.

20    The initialization process of a device will now be considered in more detail. Fig. 12 shows the main steps of the initialization process "drv\_init" of a leaf device driver (990).

25    At step 992, the driver establishes connection to the parent bus driver calling its "open" function. In calling "open", the child driver specifies call-back handlers, which will be used by the parent driver to manage the device instance  
30    driver (e.g. to shut down the driver instance). Similarly, global bus events (e.g. a catastrophic bus error) are delivered to the driver through a call-back handler.

Once the child-to-parent connection is established, the  
35    driver may use services provided by the parent driver. At step 994, the child driver asks its parent driver to make available the bus hardware resources needed for its bound device. In fact, the bus resources needed for the device have been specified as properties in the device node, and such

resources have been already allocated by the parent bus/nexus driver prior to the "drv\_init" invocation. Thus, "making available" a given bus resource (e.g. device I/O registers), means that the driver obtains an appropriate property value 5 and calls an appropriate bus/nexus service routine.

Then, at step 996, the child driver initializes the device hardware to an operational state. Finally, at step 998, the (leaf) driver performs the self-registration in the device 10 registry. This declares that a new device, together with its device driver instance, is available in the system.

Once the device (i.e. the device driver instance) is registered, a driver client may find such a device in the 15 registry (using the device registry API) and may perform operations on the device, calling driver service routines exported through the device registry entry. The driver is not obligated to register any device driver instances and offer any device service routines through the device registry. The 20 device driver instance registration is only required for clients who plan on finding their devices through the device registry. If other client-to-driver binding mechanisms are in use, the associated devices need not take part in the device registry.

25

The (leaf) "drv\_init" routine is called in the context of DKI thread. It makes possible to invoke directly the bus/nexus and DKI services allowed in the DKI thread context only.

30 Step 902 of Fig. 8 may now be described in detail. Like a leaf driver, initialization of a bus/nexus driver comprises enabling it to access its hardware (e.g. internal bus bridge registers), to establish connection with its parent bus/nexus driver and to use services implemented by its parent driver.

35

So, step 902 might include all steps 990-998 of Figure 11. However, in the preferred embodiment, when initializing a bus/nexus driver, step 998 is omitted, i.e. the nexus driver does not implement registering its bus/nexus instance and the

corresponding device in device registry 50. Thus, upon a call in operation from a driver client 250:

- the client to (leaf) driver binding, aiming at a leaf node, uses the device registry;

5 - then a binding mechanism (child to parent) is necessary to enable travelling within the tree from the leaf towards the root; rather than a child-to-parent driver binding mechanism, a parent to child action is used instead: each parent bus/nexus driver provides its child with a pointer to its 10 service routines vector and its identifier, for use in operation (and also e.g. when the "drv\_probe" or "drv\_init" routine of the child driver will be invoked). This does not use the device registry.

15 In summary, when starting the computer, a full device and driver scanning may be performed as follows. A first "Bus Driver Initialization routine" or BDIR is launched, with the "current bus/nexus node" being the (driver) root, i.e. the pci bus in the example. Loops perform a) probing, b) binding, 20 c) creation of driver instances (as appropriate in each case). This builds the child level of the current bus/nexus node. Whenever the child level comprises another bus/nexus node, phase c) launches a nested recursive call of the BDIR, with that bus/nexus node as a "current bus/nexus node".

25

Thus, considering a given nexus (sub)node:

- the child node level in the device tree is constructed by a single probing step;

30 - for the subsequent driver initialization, the subtree of the given (sub)node is explored following a branch, until a leaf-only level is found, then going back to the last nexus to follow another branch, and so on until the whole subtree has been seen.

35 Although the above described ways of scanning the device tree for i) probing, ii) binding, and iii) initializing drivers recursively are preferred, other scanning methods may be contemplated by men skilled in the art, using the same principles.

Modification of the device and driver framework in operation will now be considered. The change may be due to two different events:

- the system sees a new driver (load handler),
- 5 - a device is added/removed (hot plug).

A given nexus device driver has a handler named "Driver Load Handler", which is invoked by the parent bus/nexus driver, when a new driver appears in the system in operation: for 10 example, a new driver is downloaded at run time, and has just registered in the driver registry 40, using its "main" routine. The load handler address is given to the "given" bus/nexus driver when a connection is established between the new child driver and the parent bus/nexus driver.

15

The load handler is optional. Normally, a leaf device driver is not interested about such an event. The load handler is usually used by a bus/nexus driver (e.g. bus-to-bus bridge), which is interested in applying a newly downloaded driver to 20 its child devices which are not serviced yet.

Fig. 13 shows the actions 1010-1030 taken by the load handler. They are similar to steps 910-930 shown in fig. 8 for the bus/nexus driver initialization. The process normally 25 starts again from the root node, i.e. the pci bus in the example. Step 902 is omitted, since the bus/nexus device is already initialized.

At step 1010, the bus/nexus driver invokes the probe routines 30 registered in the driver registry matching the bus/nexus class. The probing loop of Exhibit A.7 is used again. The probe routines already invoked at initialization time will likely be invoked again. As mentioned, a probe routine must be aware about existing device nodes in order to avoid 35 creation of redundant device nodes. In addition, a probe routine must explicitly ask for bus resources in order to avoid conflicts with bus resources currently in use. In this way, the active device nodes and associated running driver

instances are protected against a perturbation caused by the run-time probing.

Thus, the probe routine implemented by the newly downloaded 5 driver will likely discover one or more device(s) previously unknown in the system, with one or more corresponding new child device nodes being created. If so, at step 1019, the bus/nexus driver will check/allocate bus resources required for such new device nodes.

10 Step 1019 is more sophisticated than corresponding step 919 of Fig. 8. In order to satisfy such a run-time resources request, the bus/nexus driver may need to confiscate resources already assigned to existing ("older") device 15 nodes; however, in so doing, the bus/nexus driver is not allowed to confiscate resources in use, i.e. resources assigned to active "older" device nodes. The driver instance associated with such an active "older" node must be shut down by the bus/nexus driver prior to any resource confiscation. 20 In order to shut down an "older" driver instance, the bus/nexus driver sends a shut down event to the "older" child driver, requesting to close the "older"-child-to-parent driver connection, i.e. to invoke "close". Once the connection is closed, the "older" device node becomes 25 inactive and the bus/nexus driver is able to confiscate bus resources assigned to that "older" node. In such a case, the bus/nexus driver should start the "older" driver instance again, invoking the driver "drv\_init" routine, to have the driver operate when its resource requirements may be 30 satisfied.

Once the bus resources allocation is done, the bus/nexus driver iterates through the child nodes and, for each inactive device node, it tries to determine a driver 35 component which should be applied to the given device, and does the binding at step 1020, which is the same as step 920. Once such a driver component is found, the bus/nexus driver calls the driver "drv\_init" routine, at step 1030, which is the same as step 930.

Finally, the bus/nexus driver invokes the child load handlers (if any) in order to propagate the loading process downstream (i.e. from parent to child). In such a way, the loading process is recursively continued by the child driver load 5 handler(s).

The load handler is called in the context of DKI thread. It makes it possible to invoke directly the bus/nexus and DKI services allowed in the DKI thread context only. 10

The case of hot-pluggable device drivers will now be considered. In case of hot-plug, the initialization/shutdown process has to be executed at runtime, rather than as part of the kernel/drivers initialization process. In such a case, 15 drivers should use DKI thread services to explicitly synchronize with already running drivers.

The hot-plug insertion event is typically reported via interrupt. In order to be notified when the hot-plug 20 insertion occurs, the bus driver connects an interrupt handler to an appropriate interrupt source.

Thus, the bus driver typically detects the insertion event at interrupt level. In order to process such an event, the bus 25 driver invokes "svDkiThreadTrigger" requesting the DKI thread to execute a bus driver routine which implements the hot-plug insertion procedure. The DKI thread serializes all actions related to the initialization and termination. In addition, it allows the insertion procedure to invoke directly the 30 parent bus/nexus and DKI services allowed in the DKI thread context only.

The hot-plug insertion processing again goes through main steps similar to steps 1010 - 1030 of figure 13, i.e.: 35

- discover a new device inserted to the bus and create the device node associated with this device,
- allocate bus resources for the new node,
- bind a driver to the new node,
- create a driver instance associated with the new node.

Actions taken by the hot-plug insertion are very close to the actions taken by the load handler, because once a new child node appears, the bus driver behaviour is always the same. It tries to allocate bus resources required for this 5 node and it tries to apply a driver to this node. It does not matter why a new device node has appeared (In case of load handler, a new device may appear because of a new probe routine having been running).

10 The driver event handler is invoked by the parent bus/nexus driver when a bus/nexus event occurs. The event handler address is given to the bus/nexus driver when a connection is established between the child driver and its parent bus/nexus driver. The event handler may be called as an interrupt 15 handler and therefore the event handler implementation must be restricted to the API allowed at interrupt level.

Among all events which are mostly bus/nexus class specific, there are three shut down related events (specified by the 20 common bus API) which are discussed in this section:

- [xxx]SYS\_SHUTDOWN - system (emergency) shut down. The prefix "xxx" may mean "DKI", "BUS", etc., to the extent the event semantics are the same for all such events. The SYS\_SHUTDOWN event notifies the driver instance that the 25 system is going to be shut down. The parent bus/nexus driver requires the child driver instance to perform an emergency shut down of the device hardware.
- DEV\_SHUTDOWN - normal device shut down. The DEV\_SHUTDOWN event notifies the driver instance that a normal device shut 30 down is requested by the bus/nexus driver.
- DEV\_REMOVAL - surprise device removal. The DEV\_REMOVAL event notifies the driver instance that the associated device has been removed from the bus/nexus and therefore the driver instance has to be shut down.

35

In the general case, the shut down event processing goes through three phases: shut down prologue, shut down mode, shut down epilogue.

The first phase (called shut down prolog) is processed synchronously within the event handler. The main purpose of the shut down prolog is to notify driver clients (or child drivers) that the shut down event has occurred, in other 5 words, to propagate the shut down prolog downstream (i.e. from parent to child). After the shut down event has been processed by driver clients (or child drivers), the driver enters in the third shut down phase (called shut down epilogue). Basically, the shut down epilogue is invoked when 10 the last reference on the driver instance goes away. Between the shut down prolog and epilogue, the driver operates in a special mode (called shut down mode). In such a mode, the driver accepts only a subset of operations from clients allowing to properly close connections to the driver.

15

Table T5 in Fig. 18 shows typical actions taken by the shut down prolog depending on the event type. Note that the notification mechanism used by the shut down prolog depends on driver type and it is different for the bus/nexus and leaf 20 (device) drivers. The shut down prolog of the bus/nexus driver notifies the child drivers calling the child driver event handlers. The shut down prolog of the leaf (device) driver notifies driver clients calling the "svDeviceEvent" routine.

25

The SYS\_SHUTDOWN prolog of a bus/nexus driver invokes the event handlers of child drivers connected to it. Once the invocation is done, the bus/nexus driver puts the hardware into a clean state. The SYS\_SHUTDOWN prolog of a leaf 30 (device) driver does not notify driver clients about the system shut down event. The driver just puts the hardware into a clean state. In fact, the SYS\_SHUTDOWN event is processed synchronously, i.e. within the event handler (In other words, the system shut down epilogue is empty). The 35 only purpose of the "SYS\_SHUTDOWN" event is to put the board hardware into a clean state in order to properly perform the system reboot (or restart).

The "DEV\_SHUTDOWN" prolog only notifies driver clients (or child drivers) that the "DEV\_SHUTDOWN" event has occurred. The real device shut down is deferred up to the "DEV\_SHUTDOWN" epilogue. The "DEV\_REMOVAL" prolog does the 5 same as the "DEV\_SHUTDOWN" one, and, in addition, aborts all I/O operations in progress because, otherwise, such operations will never be completed. Aborted operations will return to callers with an error code.

10 As soon as the shut down prolog is processed, the driver changes its internal state: it enters a shut down mode, in which it accepts only a subset of operations from clients. Basically, the driver accepts requests to abort queued operations, to release previously allocated resources, and to 15 close connection to the driver. All other operations (like open a new connection, start an I/O operation) are refused by the driver. In other words, in the shut down mode, the driver is waiting until a shut down epilogue condition is met, allowing clients to properly close existing connections to 20 the driver.

The "shut down" epilogue condition is met within a bus/nexus driver, when the last connection is closed to the driver instance. The shut down epilogue condition is met within a 25 leaf (device) driver, when the device entry is released by the last driver client, i.e. when the call-back release handler is invoked by the device registry. Note that, typically, the shut down epilogue condition is met when the driver is running in the DKI thread context. Indeed, the 30 "close" service routine and call-back release handler are both called in the DKI thread context. Thus, in natural way, the shut down epilogue is processed in the DKI thread context. It makes possible to invoke directly the bus/nexus and DKI services allowed in the DKI thread context only.

35

Table T6 in Fig. 19 shows typical actions taken by the shut down epilogue depending on the event type. The "DEV\_SHUTDOWN" epilogue puts hardware into a clean state, releases system resources used by the driver instance and, finally, closes

connection to the parent driver ("close"). The "DEV\_REMOVAL" epilogue is close to the "DEV\_SHUTDOWN" one except the device hardware is not touched by the driver, because (obviously) the device hardware is no more present on the parent bus.

5

When a shut down epilogue closes the last connection to the parent driver, the shut down epilogue condition would be met in the parent driver too. Then, the shut down epilogue is propagated upstream (i.e. from child to parent). If one of 10 driver clients (or one of child drivers) does not implement properly the shut down procedure (e.g. simply does not support the shut down), the driver may stay in the shut down mode, while never meeting the shut down epilogue condition.

15 The Hot-plug Removal will now be described.

The hot-plug removal event is typically reported via interrupt. In order to be notified when the hot-plug removal occurs, the bus driver connects an interrupt handler to an 20 appropriate interrupt source. Hot-plug removal may include:  
- "surprise removal", i.e. a device can be removed at any time with no warning. For instance, PCMCIA is a surprise removal device.  
- "non-surprise removal", i.e. the device cannot be removed 25 until the system is prepared for it. For instance, Hot-Plug CompactPCI is a non-surprise removal device.

The surprise removal interrupt notifies the bus driver that a device has been removed from the bus. The bus driver 30 interrupt handler usually detects the removed device (and associated device node) using bus specific status register(s). Once the device is detected, the interrupt handler checks whether the device node is active.

35 If the device node is inactive (i.e. there is no driver instance servicing this device), the only task of the bus driver is to update the device tree removing the device node. All bus resources associated with the node becomes free. The bus driver is not able to accomplish this task immediately at

interrupt level because the necessary services can be called in the DKI thread context only, and are not available at interrupt level. In order to satisfy the invocation context requirements, the bus driver calls "svDkiThreadTrigger" 5 requesting the DKI thread to invoke the removal procedure. Another reason to use the DKI thread is to serialize all actions related to the initialization and termination.

10 If the device node is active, the bus driver must shut down the corresponding device driver instance prior to invoke the removal procedure. In order to accomplish this task, the bus driver invokes the device driver event handler, signalling the "DEV\_REMOVAL" event. In fact, the bus driver performs the 15 shut down prolog for the given driver instance. In other words, the bus driver initiates the shut down process for the device sub-tree starting from the removed device node.

20 As the last action of the shut down event processing, the device driver will close the connection to the bus driver and, at this moment, the bus driver will perform the removal procedure. The device node removal procedure is executed in the DKI thread context because the "close" service routine is called in the DKI thread context.

25 The "non-surprise removal" interrupt requests the bus driver to enable the device removal from the bus. The processing is the same as for "Surprise Removal", except the bus driver requests the normal device shut down (DEV\_SHUTDOWN) rather than the device removal one (DEV\_REMOVAL). In addition, once 30 the device tree is updated, the bus driver enables the device removal. (The device removal enabling is usually signalled by a LED or/and by a card ejection mechanism un-locking).

35 Unloading a driver will now be discussed. The only means to unload a driver is a "Driver Unload Routine" or "drv\_unload". The "drv\_unload" is a global per driver component routine. Therefore, in order to implement unloading, the driver should handle a list of its driver instances.

The "drv\_unload" routine is optional. In case when "drv\_unload" is not provided by a driver, its "drv\_unload" field is set to NULL, and the driver code cannot be unloaded.

5 "drv\_unload" is called by the "svDriverUnregister" routine of the driver registry module, when an application wishes to unload the driver component from the system. The purpose of "drv\_unload" is to check that the driver component is not currently in use: e.g. the driver must check for each driver 10 instance whether it is locked in the device registry.

On success, all device instances are removed from the device registry and "drv\_unload" returns "K\_OK". Otherwise "K\_EBUSY" is returned, the driver component will not be unloaded, its 15 state is not changed, all registered driver instances remain in place, and the device entries are unchanged in the device registry.

20 Since the "drv\_unload" routine is called in the context of DKI thread, it can invoke directly the bus/nexus and DKI services allowed in the DKI thread context only.

25 When "drv\_unload" is called, the driver should go through the list of driver instances, and check whether any of them is currently in use. If a given instance is not used, the driver instance become invisible for potential clients. Thus, if "drv\_unload" returns "K\_OK", all previously created driver instances (if any) are deleted and all previously allocated system resources (if any) are released.

30 The driver unloading implementation is different for a bus/nexus driver, using the standard child-to-parent driver binding mechanism, and a leaf device driver, using the client-to-driver binding mechanism based on the device 35 registry. (If another client-to-driver binding mechanism is used, the driver unloading implementation is binding mechanism dependent).

The "drv\_unload" routine in a Bus/nexus Driver typically takes the following actions:

- . step 1: check that the driver component is not in use.
- "drv\_unload" iterates through the driver instances list and, 5 for each driver instance, checks whether a connection is opened to the driver instance. Once a driver instance with an open connection is found, the iteration is aborted and "K\_EBUSY" is returned. Otherwise, "drv\_unload" proceeds to step 2.
- 10 . step 2: release resources associated with the driver component. "drv\_unload" iterates through the driver instances list and, for each driver instance, releases all system resources associated with the instance and, finally, closes the connection to the parent bus. Once the iteration is 15 finished, "drv\_unload" returns "K\_OK".

The fact that "drv\_unload" is running in the DKI thread context guarantees stability of the driver instances and open connections during the "drv\_unload" execution: a new driver 20 instance may be created only by the "drv\_init" routine and a new parent-to-child connection may be open only by the "drv\_init" or "drv\_probe" routines. Both "drv\_init" and "drv\_probe" are invoked in the DKI thread context. Thus, "drv\_init" and "drv\_probe" are serialized with "drv\_unload" 25 by the DKI thread. On the other hand, if a bus/nexus driver supports hot-pluggable devices, it is up to the bus/nexus driver to implement a synchronization mechanism with a hot-plug insertion interrupt which may occur during the driver unloading.

30

The "drv\_unload" routine in a (leaf) device driver typically takes the following actions:

- . step 1: check that the driver component is not in use.
- "drv\_unload" iterates through the driver instances list and, 35 for each driver instance, invokes "svDeviceUnregister" in order to remove the driver instance entry from the device registry. If "svDeviceUnregister" fails (i.e. returns "K\_EBUSY"), the iteration is aborted and "drv\_unload"

proceeds to step 2. Otherwise (i.e. all device instances are successfully unregistered), "drv\_unload" proceeds to step 3.

· step 2 (fail): "drv\_unload" iterates through the driver instances list and, for each driver instance which has been 5 unregistered at step 1, invokes "svDeviceRegister" in order to register the driver instance again. Once the iteration is finished, and the initial state of the driver component has been restored, "drv\_unload" returns "K\_EBUSY".

· step 3: "drv\_unload" iterates through the driver instances 10 list and, for each driver instance, releases system resources associated with the instance and, finally, closes the connection to the parent bus when all resources associated with the driver component have been released. Once the iteration is finished, "drv\_unload" returns "K\_OK".

15

It should be reminded that "drv\_unload" is running in the DKI thread context. This guarantees stability of the driver instances during the "drv\_unload" execution. Indeed, a new driver instance may be created only by the "drv\_init" routine 20 which is invoked in the DKI thread context too. Thus, "drv\_init" is serialized with "drv\_unload" by the DKI thread.

As already mentioned, in some cases, a driver may defer its device initialization, until it is opened. This is a way to 25 resolve a conflict about usage of the same resource by multiple drivers. In that way, drivers sharing a resource can be loaded at same time, provided they are not opened at same time. In such a case, the initialization and shutdown process of the driver has to be executed at runtime, i.e. at time of 30 driver "open" and "close", respectively, rather than as part of the kernel/drivers initialization process. Thus, deferred initialization drivers should also use the above described DKI thread services to explicitly synchronize with already running drivers.

35

On another hand, the relation between a bus driver and its child driver(s) may offer several possibilities.

A bus driver may provide multiple API's. A typical example is a PCI bus driver providing a common bus API and a PCI bus API. The PCI bus driver API is named "pci" and specified by a "PciBusOps" structure. The common bus driver API is named 5 "bus" and specified by a "BusOps" structure, which provides a subset of services provided by "PciBusOps". Such a bus driver is able to support drivers which use either common ("bus") or PCI ("pci") parent bus interface. When the bus driver invokes a child driver, it gives a pointer to either 10 the "BusOps" or the "PciBusOps" structure depending on "bus\_class" specified in the child driver registry entry.

Conversely, a driver may be adaptive as to its parent bus API, i.e. be able to run on top of a number of different 15 buses (e.g. "pci" and "isa"). Typically, such a driver is composed of two parts: bus class specific and bus class independent. The bus class specific part of the driver code mainly deals with the device probing and initialization. In addition, it provides an abstraction layer in order to hide 20 the bus class dependencies from the bus class independent part of the driver code.

Such a multi-bus driver should be registered multiple times in the driver registry. Each entry specifies a given bus 25 class API (via the "bus\_class" field) on top of which the driver may run. In order to determine to which bus class the driver is applied, the "drv\_probe", "drv\_bind" and "drv\_init" routines have to be entry specific (i.e. bus class API specific). Under such conditions, when a given driver's 30 routine is invoked by a bus driver, the driver detects the bus class to which it is applied and casts the "bus\_ops" argument to the appropriate structure (e.g. "PciBusOps").

In this invention, as it will be appreciated, each driver has 35 driver management functions and driver/device service functions. In accordance with the device tree (and the corresponding virtual driver tree), the functions are closely related to the needs of the driver to device relation. This is advisable to avoid that shutdown of a driver instance be

detimental to operation of other <device/driver instance> pairs.

Thus, the described system provides:

- 5 - in operation, a simple tool for enabling driver client code to perform operations on a given device, using the device registry;
- a flexible and consistent way of initializing drivers and devices, when starting the computer; //
- 10 - a flexible and consistent way of taking new drivers into account, as well as to unload drivers, in operation;
- a flexible and consistent way of managing device insertion and/or removal (and/or failure), in operation;
- a consistent way of managing resources shortage, when appropriate.

These features may be used in any desired combination. In such combinations, the features being useless may be omitted.

- 20 This invention is not restricted to the embodiment as disclosed. Although it extensively refers to ChorusOS, it may be applied to other operating systems as well. The organization of the probing, binding and initialization operations and loops may be different. The organization of data and code may
- 25 also be different; this remark concerns the drivers and driver instances, the device registry, the driver registry, the device tree, and their relations. Also, despite the driver tree is described as virtual, it may also be given a data representation.

## EXHIBIT A

## A-1

```
typedef struct DevRegEntry {  
    5    char* dev_class; /* device class */  
        void* dev_ops; /* device driver service routines */  
        void* dev_id; /* device driver ID */  
        DevNode dev_node; /* device node */  
    } DevRegEntry;
```

10

## A-2.1

```
void  
svDkiThreadCall (DkiCallToken* token, DkiCallHandler handler,  
    void* cookie);
```

15

## A-2.2

```
void  
svDkiThreadTrigger (DkiCallToken* token, DkiCallHandler handler,  
    void* cookie);
```

20

## A-3.1

```
DevNode  
dtreeNodeRoot ();  
DevNode  
25    dtreeNodeChild (DevNode node);  
DevNode  
dtreeNodePeer (DevNode node);  
DevNode  
dtreeNodeParent (DevNode node);
```

30

## A-3.2

```
DevNode  
dtreeNodeAlloc ();
```

```
void
dtreeNodeFree (DevNode);
void
dtreeNodeAttach (DevNode pnode,DevNode cnode);
5 void
dtreeNodeDetach (DevNode node);
```

#### A-3.3

##### DevProperty

```
10 dtreePropFind (DevNode node, char* name);
DevProperty
dtreePropFindNext (DevProperty prop, char* name);
unsigned int
dtreePropLength (DevProperty prop);
15 void*
dtreePropValue (DevProperty prop);
char*
dtreePropName (DevProperty prop);
DevProperty
20 dtreePropAlloc (char* name)
int length);
void
dtreePropFree (DevProperty prop);
void
25 dtreePropAttach (DevNode node, DevProperty prop);
void
dtreePropDetach (DevProperty prop);
```

#### A-3.4

##### DevNode

```
dtreeNodeAdd (DevNode parent, char* name);
DevNode
dtreeNodeFind (DevNode parent, char* name);
```

```
DevProperty
dTreePropAdd (DevNode node,
char* name,
void* value,
5     unsigned int length);
```

## A-4.1

```
svDeviceAlloc
typedef void (*DevRelHandler) (DevRegEntry*);
10    DevRegId
        svDeviceAlloc (DevRegEntry* entry,
                     unsigned int version
                     Bool shared,
                     DevRelHandler handler);
```

15

```
A-4.2
svDeviceRegister
void
svDeviceRegister (DevRegId dev_id);
```

20

```
A-4.3
svDeviceUnregister
KnError
svDeviceUnregister (DevRegId dev_id);
```

25

```
A-4.4
svDeviceEvent
typedef enum {
    DEV_EVENT_NULL = 0, /* NULL (unused) event */
30    DEV_EVENT_SHUTDOWN, /* normal device shut down request */
    DEV_EVENT_REMOVAL /* hot-plug device removal notification */
} DevEvent;
void
```

```
svDeviceEvent (DevRegId dev_id, DevEvent event,  
void* arg);
```

A-4.5

5 svDeviceFree

void

```
svDeviceFree (DevRegId dev_id);
```

A-4.6

10 svDeviceLookup

```
typedef void (*DevEventHandler) (void*, DevEvent, void*);
```

KnError

```
svDeviceLookup (char* dev_class,
```

unsigned int dev\_version,

15 unsigned int dev\_unit,

DevEventHandler cli\_handler,

void\* cli\_cookie,

DevClientId\* cli\_id);

20 A-4.7

svDeviceEntry

DevRegEntry\*

```
svDeviceEntry (DevClientId* cli_id);
```

25 A-4.8

svDeviceRelease

void

```
svDeviceRelease (DevClientId* cli_id);
```

30 A-5

```
typedef struct DrvRegEntry {  
char* drv_name; /* driver name */  
char* drv_info; /* driver info */
```

```
char* bus_class; /* class of the parent API (e.g. "pci") */
int bus_version; /* minimal version of the parent class API */
/*
* drv_probe() performs an enumeration/probing of
5  * device(s) residing on the parent bus. When such a device is found,
* the probe routine creates a new child device node (with respect to
* the parent "bus_node") in the device tree. The probe routine
* attaches device properties (including required resources)
* to the newly created device node.
10 */
void
(*drv_probe) (DevNode bus_node,
void* bus_ops,
void* bus_handle);

15 /*
* drv_bind() performs a driver-to-device binding.
* It gives an opportunity to the driver to dynamically
* bind the driver to a given device. The device is specified
* by its device node in the device tree.
20 */
void
(*drv_bind) (DevNode dev_node);

/*
* drv_init() creates an instance of driver for
25 * the given device. Device is specified by its device
* node in the device tree.
*/
void
(*drv_init) (DevNode dev_node,
30 void* bus_ops,
void* bus_handle);

/*
* drv_unload() notifies the driver that somebody wishes to
```

\* unload the driver component from the system.  
\* The driver checks whether it is currently used by a driver client.  
\* The K\_EBUSY return value means that the driver code is currently  
\* in use.

5    \*/

KnError  
(\*drv\_unload) ()  
} DrvRegEntry;

10    A-6.0

svDriverRegister  
KnError  
svDriverRegister (DrvRegEntry\* drv\_entry);

15    A-6.1

svDriverLookupFirst  
DrvRegId  
svDriverLookupFirst();

20    A-6.2

svDriverLookupNext  
DrvRegId  
svDriverLookupNext (DrvRegId drv\_id);

25    A-6.3

svDriverRelease  
void  
svDriverRelease (DrvRegId drv\_id);

30    A-6.4

svDriverEntry  
DrvRegEntry\*  
svDriverEntry (DrvRegId drv\_id);

## A-6.5

```
svDriverCap
KnCap*
svDriverCap (DrvRegId drv_id);
```

5

## A-6.6

```
svDriverUnregister
KnError
svDriverUnregister (DrvRegId drv_id);
```

10

## A.7

```
DrvRegId drv_curr;
drv_curr = svDriverLookupFirst();
while (drv_curr) {
15    DrvRegId drv_prev;
    DrvRegEntry* drv_entry;

    drv_entry = svDriverEntry(drv_curr);

20    if (!strcmp(my_class, drv_entry->bus_class) &&
        (my_version >= drv_entry->bus_version) &&
        drv_entry->drv_probe) {
        drv_entry->drv_probe(my_node, my_ops, my_id);
    }
25    drv_prev = drv_curr;
    drv_curr = svDriverLookupNext(drv_prev);
    svDriverRelease(drv_prev);
}
```

30

## A.8

```
DrvRegId drv_curr;
```

```

drv_curr = svDriverLookupFirst();
while (drv_curr) {
    DrvRegId drv_prev;
    DrvRegEntry* entry;
5     entry = svDriverEntry(drv_curr);
    if (entry->drv_bind && !strcmp(my_class, entry->bus_class) &&
        (my_version >= entry->bus_version)) {
        DevNode dev_node = dtreeNodeChild(my_node);
        while (dev_node) {
10       entry->dev_bind(dev_node);
        dev_node = dtreeNodePeer(dev_node);
        }
    }
    drv_prev = drv_curr;
15     drv_curr = svDriverLookupNext(drv_curr);
    svDriverRelease(drv_prev);
}

```

## A.9

```

20     DevNode dev_node;
    dev_node = dtreeNodeChild(my_node);
    while (dev_node) {
        if (!dtreePropFind(dev_node, "active")) {
            DevProperty drv_prop;
25         char* drv_name;
            DrvRegId drv_curr;
            drv_prop = dtreePropFind(dev_node, "driver");
            if (drv_prop) {
                drv_name = (char*) dtreePropValue(drv_prop);
30         drv_curr = svDriverLookupFirst();
                while (drv_curr) {
                    DrvRegId drv_prev;
                    DrvRegEntry* drv_entry;

```

```
drv_entry = svDriverEntry(drv_curr);
if (!strcmp(my_class, drv_entry->bus_class) &&
    (my_version >= drv_entry->bus_version) &&
    !strcmp(drv_name, drv_entry->drv_name) &&
    drv_entry->drv_init) {
    5    drv_entry->drv_init(dev_node, my_ops, my_id);
    if (dtreePropFind(dev_node, "active")) {
        svDriverRelease(drv_curr);
        break;
    }
    10   }
    }
    drv_prev = drv_curr;
    drv_curr = svDriveLookupNext(drv_prev);
    svDriverRelease(drv_prev);
    15   }
    }
    dev_node = dtreeNodePeer(dev_node);
}
20
```

## Exhibit B

## B-1 - "DevRegEntry" structure

5 The fields of the "DevRegEntry" structure are as follows:

- "dev\_class" points to a string specifying the device class name.
- "dev\_ops" points to a structure of driver service routines.
- "dev\_id" points to a handle which is usually passed back to the driver (as first argument) each time a driver service routine is invoked.

10 - "dev\_node" specifies the device node in the tree which is serviced by driver. Note that a driver client must switch to the DKI thread context when accessing the device tree structure.

- The "version" argument specifies the driver interface version implemented by the device driver instance. Basically, the "dev\_class" field and "version" argument specify

15 API implemented by the driver. Such an API is provided to driver clients as a structure of indirect functions implementing the API service routines. For instance, the "uart" device API is specified by the "UartDevOps" structure, the "timer" device API is specified by the "TimerDevOps" structure. A pointer to such a structure is obtained by a device driver client from the "dev\_ops" field of the device registry entry.

20 - The "shared" argument specifies if the device driver instance may be shared between multiple clients. In other words, it indicates if the registered instance may be looked up multiple times.

- The "handler" argument specifies the driver handler which is invoked by the device registry module as an acknowledge to a shut-down event (see "svDeviceEvent").

25 "handler" is called when the last reference to the device registry entry goes away and a shut-down event has been previously signalled on the entry. "entry" is passed back to the driver as the argument of "handler". In case of success, a non zero "DevRegId" is returned, otherwise NULL is returned.

30 Device registry assigns a logical unit number to the physical device. The logical unit number is unique within "dev\_class". Basically, the device registry handles a logical unit counter per device class. All counters are initialized to zero. Each time a new device entry is allocated, the current counter value is assigned to the entry logical unit number

and the counter is incremented. Thus, the logical device order within class corresponds to the allocation order. The logical unit number is mainly used by device driver clients to iterate through a given device class looking for a certain device. Using a valid logical unit number, a client is able to access the corresponding device registry entry, in 5 particular, the "dev\_node" field which points to the device node. The physical device corresponding to a given logical unit number may then be detected by the device node properties and the node position within the device tree. Note that the device driver may allocate (and then register) multiple entries in order to declare different classes (APIs) for the same physical device. For instance, a driver may implement orthogonal sets of 10 operations like power management which are additional to the main device functionality. Another example is a device driver servicing a multi-function device. Such a driver would register multiple entries: one entry per device function. Note that, from clients point of view, such a device would be visible as multiple (mono-function) independent devices.

15

## B-2 - Basic calls of DKI Thread

B-2.1. svDkiThreadCall (A-2.1) synchronously invokes a routine in the context of the DKI thread. Synchronously means that the caller is blocked until the invoked routine 20 returns.

B-2.2. svDkiThreadTrigger (A-2.2), asynchronously invokes a routine in the context of the DKI thread. Asynchronously means that the function immediately returns to the caller without waiting for the invoked routine to return. If the driver needs to know 25 when the handler returns, it should use any synchronization mechanism inside the handler itself.

In both case, the "token" argument is the address of a DkiCallToken structure which must be allocated by the caller. This structure is opaque for the driver, and is used only 30 by the DKI thread. The same structure may be reused for subsequent call (once the handler is invoked, in the case of "svDkiThreadTrigger"). The "handler" argument specifies the routine to call, and the "cookie" argument specifies the argument to pass to the "handler" routine when called.

### B-3 - Device tree functions

#### B-3.1 - Tree browsing

5

"dtreeNodeRoot" returns the root device node if the device tree is not empty, otherwise NULL is returned.

10 "dtreeNodeChild" returns the first child node from the children list. NULL is returned when the children list is empty.

"dtreeNodePeer" returns the next device node from the sibling list, if any, otherwise NULL is returned. The "node" argument specifies the current device node in the sibling list.

15

"dtreeNodeParent" returns the parent device node, if any, otherwise NULL is returned. The "node" argument specifies the child device node.

#### B-3.2 - Node management

20

"dtreeNodeAlloc" allocates a new device node object. A non zero "DevNode" cookie is returned in case of success, otherwise NULL is returned. Allocated node has neither parent nor child nodes. There are no properties attached to the newly allocated node.

25

"dtreeNodeFree" releases the memory allocated by the node object and all property objects attached to the node. The "node" argument specifies the node object being released.

30

"dtreeNodeAttach" adds the device node specified by the "cnode" argument in the child list of the parent node specified by the "pnode" argument.

"dtreeNodeDetach" detaches the node object specified by the "node" argument from its parent (if any). When "node" specifies the root node, the device tree becomes empty.

### B-3.3 - Node properties

"DevProperty" is an abstract type designating a device property object, and is opaque for the driver.

5

"dtreePropFind" searches the first property within the properties list of the device node specified by the "node" argument. If the "name" argument is not NULL, "dtreePropFind" returns the first property whose name matches the "name" string. If the "name" argument is NULL, "dtreePropFind" returns the first property from the list regardless of its name. In case of success, a non zero "DevProperty" cookie is returned, otherwise NULL is returned. Once the first property is found, the "dtreePropFindNext" routine may be used in order to find a subsequent property in the list.

10

"dtreePropFindNext" searches the next property within the properties list. The current position within the properties list is specified by the "prop" argument. If the "name" argument is not NULL, "dtreePropFindNext" returns the next property (with respect to the current position) whose name matches the "name" string. If the "name" argument is NULL, "dtreePropFindNext" returns the next property from the list (with respect to the current position) regardless of its name. In case of success, a non zero "DevProperty" cookie is returned, otherwise NULL is returned. The "dtreePropFind" and "dtreePropFindNext" pair is typically used to iterate through either all device node properties or a subset of device node properties whose names match a given name. In case when an user knows that there is only one property with a given name attached to the device node, a single "dtreePropFind" invocation is enough.

15

"dtreePropLength" returns the property value length (in bytes). The property object is specified by the "prop" argument.

20

"dtreePropValue" returns a pointer to the first byte of the property value. The property object is specified by the "prop" argument. A driver can read and write the property value directly using the returned pointer. Typically, the driver will cast the pointer to a well known type/structure in order to access the property value. Note that the property value must always be presented in the CPU endian format.

"dtreePropName" returns a pointer to an ASCII string which designates the property name. The property object is specified by the "prop" argument.

5 "dtreePropAlloc" allocates a new device property object. A non zero "DevProperty" cookie is returned in case of success, otherwise NULL is returned. The "name" argument specifies the property name. The "length" argument specifies length of the property value. The property value is undefined. The allocated property object is not attached to any device node. Once the property value is initialized, the property object can be attached to a node (i.e. added to the node properties list) using the 10 "dtreePropAttach" function.

15 "dtreePropFree" releases the memory allocated by the property object. The "prop" parameter specifies the property object being released. Note that the property object must not be attached to any device node.

20 "dtreePropAttach" adds the property object specified by the "prop" argument to the node properties list. The "node" argument specifies the device node to which the property will be attached.

25 "dtreePropDetach" detaches the property object from the device node to which it is attached (if any). The "prop" argument specifies the property object.

#### B-3.4 - high level services

30 "dtreeNodeAdd" allocates a new device node object, and add it to the children list of the given "parent" device node. The "name" argument specifies the name of the new allocated device node. That means a "node" property is allocated and attached to the new node having the value specified by "name". In case of success, the new allocated node object is returned, otherwise, a NULL pointer is returned.

35 "dtreeNodeFind" looks for a named node in the children list of a given device node. The "parent" argument specifies the device node of which the children list should be searched. The "name" argument specifies the value which must match the "node"

property value of the searched node. In case of success, the matching node object is returned, otherwise a NULL pointer is returned.

"dtreePropAdd" allocates a new property, sets its value and attaches it to a given device node. The "node" argument specifies the node to which the new property must be attached. The "name" argument specifies the name of the new property to allocate. The "length" argument specifies the memory size in bytes to allocate for the new property value. The "value" argument specifies the value to be set for the new allocated property. In case of success, the new allocated property object is returned, otherwise, 10 a NULL pointer is returned.

#### B-4 - Device registry operations

B-4.1. "svDeviceAlloc" in A-4.1 allocates a device registry entry for a given device driver instance. Note that the entry is allocated in invalid state. This means that the entry is not included in the registry, i.e. it is not visible for clients via "svDeviceLookup". The entry becomes valid (i.e. visible for clients) when a subsequent "svDeviceRegister" is invoked. On the other hand, the driver is allowed to call "svDeviceEvent" on such an entry. Shut down events signalled on an invalid entry are memorized by the device 20 registry and they are processed when the entry becomes valid (i.e. when "svDeviceRegister" is called). The "entry" argument points to the "DevRegEntry" structure which designates the device driver instance.

B-4.2. "svDeviceRegister" in A-4.2 adds a given device registry entry to the registry. 25 The entry must be previously allocated by "svDeviceAlloc". The entry becomes valid (and therefore visible for clients) only if there is no shut-down event signalled on the entry. Otherwise, the entry remains invalid and the device registry module invokes "handler" previously specified via "svDeviceAlloc". A pointer to the "DevRegEntry" structure (previously specified via "svDeviceAlloc") is passed back to the driver as the 30 argument of "handler". The "handler" routine is called in the DKI thread context. Note that when "handler" is called the entry is no more registered. Thus, it is useless to invoke "svDeviceUnregister" once "handler" is called. On the other hand, in order to free memory resources, the driver should release the device registry entry invoking

"svDeviceFree" once the entry is no more used by the driver (Typically, the entry is no more used by the driver when the connection to the parent bus/nexus driver is closed, i.e. when the parent driver is unable to invoke a call-back handler signalling an bus/nexus event).

5

B-4.3. "svDeviceUnregister" in A-4.3 removes the device entry specified by the "dev\_id" argument from the device registry (if the entry is valid). "svDeviceUnregister" returns the following results:

- "K\_OK" if the device entry is successfully removed from the registry.
- 10 - "K\_EBUSY"
  - . if the device entry is not removed from the registry because it is locked by a driver client (the "svDeviceLookup/Release" pairs are not balanced), or
  - . when an invalid entry is specified.

The driver must remove all related device entries from the registry prior to be unloaded  
15 from the system. Typically, "svDeviceUnregister" is issued by the device driver (for each device driver instance) when the driver code is requested to be unloaded via "drv\_unload". The driver unload fails when the driver code is still used by the system.

B-4.4. "svDeviceEvent" in A-4.4 notifies the device registry module that a shutdown  
20 event has occurred (e.g. hot-pluggable device has been removed). The shutdown event basically means that the driver instance must not be used any more by driver clients and the device entry must be removed from the registry (i.e. the driver instance will disappear). The shutdown event is immediately propagated to all driver clients in order to stop the device usage and to release the device entry as soon as possible. The device  
25 is specified by the "dev\_id" argument.

"svDeviceEvent" propagates the shutdown event invoking the "event handler" routines specified by the driver clients in "svDeviceLookup". The "event" and "arg" arguments are opaque for the device registry. They are passed (as arguments) to the client event handlers. "event" specifies the reason of shut-down as follows:

- 30 - DEV\_EVENT\_SHUTDOWN for normal device shut down,
- DEV\_EVENT\_REMOVAL for hot-plug (surprise) device removal.

"arg" is "event" specific. All driver clients are requested to release the device entry (as soon as possible) invoking the "svDeviceRelease" routine. Note that, prior to the shutdown event propagation, "svDeviceEvent" removes the device entry from the registry in order to prevent the entry to be found (and locked) by new driver clients.

5 Once all driver client handlers are invoked, "svDeviceEvent" returns to the driver. Note that the real device shutdown has to be deferred until the "handler" routine invocation.

Once the device entry is released by the last driver client, the device registry module invokes "handler" previously specified via "svDeviceAlloc". A pointer to the

10 "DevRegEntry" structure (previously specified via "svDeviceAlloc") is passed back to the driver as the argument of "handler". The "handler" routine is called in the DKI thread context. Typically, "svDeviceEvent" is used by a device driver servicing a hot-pluggable device.

15 "svDeviceEvent" is invoked by the device driver when the driver is notified (by its parent driver) that the device is removed from the bus. Note that in case of hot-plug removal, the device driver must be still operational until the device entry is released (ignoring all requests to the driver, for example). Note that the driver is allowed to call "svDeviceEvent" on an invalid (unregistered) entry. In such a case, "handler" will be

20 invoked only when an attempt of registration will be made, i.e. "svDeviceRegister" will be invoked. In addition, such a registration will fail and the entry will remain invalid.

B-4.5. "svDeviceFree" in A-4.5 releases a given device registry entry previously allocated by "svDeviceAlloc". The device entry being released is specified by the

25 "dev\_id" argument. The entry being released must be invalid (unregistered). Thus, if the entry has been registered by "svDeviceRegister", the driver is allowed to release it if either the entry is successfully unregistered by "svDeviceUnregister" or "handler" (previously specified by "svDeviceAlloc") is invoked.

30 B-4.6. "svDeviceLookup" in A-4.6 searches the device entry in the registry matching the specified device class and logical unit. The arguments are as follows:

- "dev\_class" specifies the device class.
- "dev\_version" specifies the minimum device driver interface version required.

- "dev\_unit" specifies the logical device unit in the class.
- "cli\_handler" specifies the event handler which is called when a device event is signalled.
- "cli\_cookie" specifies the first argument of "cli\_handler".

5 - "cli\_id" is an output argument identifying the client token on the matching device entry. This "cli\_id" is passed back as an argument to "svDeviceEntry" to get the associated device entry, or to "svDeviceRelease()" to release the device driver instance.

"svDeviceLookup" returns the following results:

- "K\_OK" if the device entry matching the "dev\_class" and "dev\_unit" pair is found in the registry.
- "K\_EBUSY" if the device entry matching the "dev\_class" and "dev\_unit" pair is found in the registry. On the other hand, the device is already looked up by another client and the device sharing is not allowed.
- "K\_EUNKNOWN" if there is no device entry in the registry matching "dev\_class".

15 - "K\_UNDEF" if there is no device entry in the registry matching "dev\_unit" in the "dev\_class". On the other hand, there is a device entry in the "dev\_class" with logical unit greater than "dev\_unit". In other words, "dev\_unit" falls into a hole in the "dev\_class" unit ranges

- or if the device entry matching the "dev\_class" and "dev\_unit" pair found in the 20 registry, implements a version of interface older than the required one.
- "KETOOMUCH" if there is no device entry in the registry matching "dev\_unit" in the "dev\_class". Moreover "dev\_unit" is greater than all existing logical units in the "dev\_class".

25 In case of success, the corresponding device entry is locked in the registry until a subsequent "svDeviceRelease". Note that the device registry lock may or may not be exclusive, depending on the value of the "shared" argument specified in "svDeviceAlloc". In other words, the device may be shared between multiple driver clients, if "shared" was TRUE at registration time.

30

The device driver instance must not disappear while its device entry is locked in the registry. In case of a catastrophic device shutdown (e.g. hot-plug removal) the device

entry lock owners are notified (through the "cli\_handler" routine invocation) that the device entry must be released as soon as possible.

5 B-4.7. "svDeviceEntry" in A-4.7 returns the device entry associated with a given client identifier. The "cli\_id" argument specifies the client identifier previously returned by "svDeviceLookup".

10 B-4.8. "svDeviceRelease" in A-4.8 routine releases the lock on the given device entry. The device entry is specified by the "cli\_id" argument. Obviously, a device driver client must no more access the device driver instance and the device node once the device entry has been released.

#### B-5 - "DrvRegEntry" structure

15

The "DrvRegEntry" structure has the following fields:

- "drv\_name" points to a string specifying the driver name (e.g. driver file name).
- "drv\_info" points to a string specifying an extra information about driver component (e.g. version, author).
- 20 - "bus\_class" points to a string specifying the class of the parent driver API required for the driver (e.g. "pci").
- "bus\_version" specifies the minimal version of the parent driver API required for the driver. Note that if a bus/nexus driver provides an API version which is less than "bus\_version", the corresponding driver component will never be called by such a bus/nexus driver. In other words, the "drv\_probe", "drv\_bind" and "drv\_init" routines will never be invoked by such a bus/nexus driver.
- 25 - "drv\_probe" points to a static driver routine which performs the device enumeration/probing on the bus.
- "drv\_bind" points to a static driver routine which performs the driver-to-device binding.
- 30 - "drv\_init" points to a static driver routine which clones an instance of driver for the given device.

- "drv\_unload" points to a static driver routine which is invoked by the driver registry module when somebody wishes to unload the driver code from the system.

## B-6 - Driver registry operations

5

B-6.0. "svDriverRegister" adds the driver entry to the driver registry. It returns "K\_OK" in case of success, otherwise K\_ENOMEM is returned to indicate the system is out of memory. The "drv\_entry" argument points to a "DrvRegEntry" structure specifying driver properties and static driver routines.

10

B-6.1. "svDriverLookupFirst" in A-6.1 returns the first driver entry in the registry. When the registry is not empty a non zero "DrvRegId" (designating the entry) is returned, otherwise NULL is returned. In case of success, the driver entry is locked in the registry. It should be unlocked by a subsequent invocation of "svDriverRelease" or "svDriverUnregister". The driver entry lock is not exclusive. It is just a counter which prevents the driver component to be unloaded when it is still used by a bus/nexus driver for the probing or initialization.

20

B-6.2. "svDriverLookupNext" in A-6.2 returns the next driver entry in the registry. The current entry is specified by the "drv\_id" argument. The current entry must be locked by a previously called "svDriverLookupFirst" or "svDriverLookupNext". If the current entry is not the last one in the registry a non zero "DrvRegId" (designating the entry) is returned, otherwise NULL is returned. In case of success, the next driver entry is locked in the registry. It should be unlocked by a subsequent invocation of "svDriverRelease" or "svDriverUnregister".

B-6.3. "svDriverRelease" in A-6.3 releases the lock of the driver entry specified by the "drv\_id" argument.

30

B-6.4. "svDriverEntry" in A-6.4 returns a pointer to the driver entry structure specified by the "drv\_id" argument. The driver entry being accessed must be previously locked using "svDriverLookupFirst" or "svDriverLookupNext". Note that the driver entry structure is read-only.

B-6.5. "svDriverCap" in A-6.5 returns a pointer to the driver actor capability. The driver entry is specified by the "drv\_id" argument. The driver entry being accessed must be previously locked using "svDriverLookupFirst" or "svDriverLookupNext". Note that if a given driver entry is registered by a builtin driver, a NULL pointer is returned. In 5 other words, the driver actor capability makes sense only for dynamically loaded drivers. The actor capability may be used by an application in order to delete the driver actor once the driver entry is unregistered. Note that the driver capability structure is read-only.

10 B-6.6. "svDriverUnregister" in A-6.6 tries to remove the driver entry specified by the "dev\_id" argument from the driver registry. The device entry being removed must be previously locked using "svDriverLookupFirst" or "svDriverLookupNext". In case of success, "K\_OK" is returned, otherwise "K\_EBUSY" is returned. Basically, the "K\_EBUSY" result means that either the driver entry is locked in the driver registry (i.e. 15 a static driver routine is currently used by a bus/nexus driver) or an instance of device driver is locked in the device registry (i.e. there is a driver instance which is currently used by a driver client). When "K\_EBUSY" is returned, the driver entry remains locked in the registry and should be unlocked explicitly by "svDriverRelease".

20 B.7

The probe routine gives an opportunity to the driver to discover a device (which can be serviced by the driver) residing on the bus/nexus and to create the device node associated with this device in the device tree.

25 B.8

The "drv\_bind" routine is invoked by a bus/nexus driver when "bus\_class" specified in the registry entry matches the bus/nexus driver class. "drv\_bind" is called with one argument ("dev\_node") specifying a given device node. The "drv\_bind" routine gives an opportunity to the driver to perform a driver-to-device binding. Typically, the driver 30 examines properties attached to the device node in order to determine the type of device and to check whether the device may be serviced by the driver. (The properties taken into account by the driver are typically bus architecture specific. For instance, a PCI driver would examine the vendor and device identifier properties.).

If the check is positive, the driver binds the driver to the device node attaching a "driver" property to the device node. The property value specifies the driver name. The parent bus driver uses such a property to determine the name of driver servicing the device. In other words, via the "driver" property, the child driver gives its name to the 5 parent bus driver asking to invoke the "drv\_init" routine on that device. Note that the "drv\_bind" routine should normally do nothing if a "driver" property is already present in the device node. In other words, "drv\_bind" should not override an existing (most likely explicit) driver-to-device binding. The "drv\_bind" routine is optional. In case when the "drv\_bind" routine is not provided by the driver, the "drv\_bind" field must be 10 set to NULL.

### B.9

Once the driver-to-device binding process is finished, the bus/nexus driver iterates through the child nodes and, for each device node, it tries to determine a driver 15 component which should be applied to the given device. Once such a driver component is found, its "drv\_init" routine is invoked the bus/nexus driver.

The "active" property is used to filter out already active device. In addition, once the "drv\_init" routine returns, the bus/nexus driver checks whether the device became active, and if so, it terminates iteration for that device node.

20 If the child device is not a leaf one, the initialization process is recursively continued by the "drv\_init" routine of the child driver.

Claims

1. A computer comprising:
  - a processor (110), a memory (115), and devices (118-138),
  - 5 - a program memory area, for storing an operating system (205), drivers (218-238), and driver client code (250),
  - each driver having a global driver section (6) comprising driver program code (60) and global driver data (61), both defining driver operations associated with a corresponding
  - 10 class of devices,
  - for operative ones of the devices, an associated driver instantiation (71) comprising local driver data, having an internal pointer to the global driver section of the instantiated driver,
  - 15 - the local driver data in at least certain of the driver instantiations further having at least one external pointer to another driver instantiation,
  - a device registry (50), forming a list of device registry entries (711), each corresponding to a respective one of the
  - 20 driver instantiations,
  - said driver client code (250) being capable of performing operations on a given device by:
    - . fetching in said device registry (50) a pointer to the driver instantiation (71) associated with the given device,
    - 25 . performing the operations on the given device by using operations provided by the associated driver instantiation, directly or using further operations reached in other driver instantiations by using iteratively said external pointer.
- 30 2. A computer according to claim 1, wherein said local driver data further comprise local device data, defining a current condition of the device associated with the driver.
- 35 3. A computer according to claim 1, further comprising device tree data (500-536), defining a tree representation of at least some of said devices, in which each device node represents a device and its device characteristics, and identifies at most one driver capable of managing that device.

4. A computer according to claim 3, wherein said device registry entry (711) further comprises a node identifier (715), defining an associated node in the device tree.
5. A computer according to claim 3, wherein said global driver section comprises code (60) adapted to implement at least one of the following special operations:
  - probing driver serviceable devices,
  - binding a driver serviceable device with the driver,
  - initializing a driver instantiation with the bound device,
  - unloading an existing driver instantiation.
6. A computer according to claim 5, further comprising a driver registry (40), forming a list of driver registry entries, each comprising at least one pointer to said special operations.
7. A computer according to claim 6, wherein said driver registry (40) directly or indirectly defines device-independent driver characteristics, comprising a driver identifier and a parent driver class identifier.
8. A computer according to claim 6, wherein said global driver section (6) comprises code adapted to self register the driver in said driver registry.
9. A computer according to claim 6, wherein, for at least intermediary ones of the drivers, said special operation of initializing the driver further comprises sub-node initialization operations (910-930).
10. A computer according to claim 9, wherein said sub-node initialization operations comprise inspecting the driver registry to find child drivers which match the driver being considered (913, 923).
11. A computer according to claim 10, wherein said sub-node initialization operations comprise initiating a special

operation of probing driver serviceable devices in such child drivers (918).

12. A computer according to claim 6 or claim 11, wherein said special operation of probing driver serviceable devices (918) comprises creating corresponding device nodes in the device tree.

13. A computer according to claim 12, wherein said sub-node initialization operations further comprise allocating resources to each child driver serviceable device (919).

14. A computer according to claim 12, wherein said sub-node initialization operations further comprise binding a child driver with a corresponding device in said device tree (920).

15. A computer according to claim 14, wherein said sub-node initialization special operations further comprise initializing the child driver (930), which results into a nested implementation (940) of said sub-node initialization operations when said child driver is itself an intermediate driver.

16. A computer according to claim 5, wherein, for at least other drivers than the intermediary ones, said driver special operation of initializing a driver instantiation for the bound device further comprises registering the new driver instantiation in the device registry (998).

17. A computer according to claim 5, wherein said operating system contains code for launching said sub-node operations in a driver bound to a given device in the device tree (207).

18. A computer according to claim 17, wherein said code (207) for launching said sub-node operations comprises a synchronizing thread in the nucleus of said operating system.

19. A computer according to claim 17, wherein said operating system is active upon computer start to initiate said code

(207) for launching said sub-node operations in a driver bound to the main bus.

20. A computer according to claim 17, wherein said operating system is active upon detection of a new driver to initiate said code (207) for launching said sub-node operations in a driver bound to the main bus.

5 21. A computer according to claim 17, wherein said operating system is active, upon detection of a new device by an intermediate device, to initiate said code (207) for launching said sub-node operations in the driver corresponding to the intermediate device.

10 22. A computer according to claim 17, wherein said operating system is arranged to de-activate the corresponding driver instantiation, upon detection of a device removal.

15 23. A method of managing drivers in a computer, comprising:  
20 a. providing each driver with a global driver section (6) comprising driver program code (60) and global driver data (61), both defining driver operations associated with a corresponding class of devices,

25 b. for each operative one of the devices, providing an associated driver instantiation (71) comprising local driver data, having an internal pointer to the instantiated driver,  
c. providing a device registry (50), forming a list of device registry entries, each corresponding to a respective one of the driver instantiations,

30 whereby a device call may be directed to the device registry, to reach the driver instantiation corresponding to the device being called.

35 24. The method of claim 23, wherein said local driver data of step b. further comprise local device data, defining a current condition of the device associated with the driver.

25. The method of claim 23, further comprising the step of:

d. providing device tree data (500-536), defining a tree representation of at least some of said devices, in which each device node represents a device and its device characteristics, and identifies at most one driver capable of 5 managing that device.

26. The method of claim 25, wherein each said device registry entry (711) of step c. further comprises a node identifier, defining an associated node in the device tree.

10 27. The method of claim 25, wherein said global driver section (6) of step a. further comprises code adapted to implement at least one of the following special operations:  
a1. probing driver serviceable devices,  
15 a2. binding a driver serviceable device with the driver,  
a3. initializing a driver instantiation with the bound device,  
a4. unloading an existing driver instantiation.

20 28. The method of claim 27, further comprising the step of:  
e. providing a driver registry (40), forming a list of driver registry entries, each comprising at least one pointer to said special operations.

25 29. The method of claim 28, wherein said driver registry (40) of step e. directly or indirectly defines device-independent driver characteristics, comprising a driver identifier and a parent driver class identifier.

30 30. The method of claim 29, wherein said global driver section (6) of step a. comprises code adapted to self register the driver in said driver registry.

35 31. The method of claim 29, wherein, for at least intermediary ones of the drivers, said special operation a1. of initializing the driver further comprises sub-node initialization operations (910-930).

32. The method of claim 31, wherein said sub-node initialization operations comprise:

a10. inspecting the driver registry to find child drivers which match the driver being considered (913, 923).

5

33. The method of claim 32, wherein said sub-node initialization operations comprise

a11. initiating a special operation of probing driver serviceable devices in such child drivers (910).

10

34. The method of claim 27 or claim 32, wherein said special operation of probing driver serviceable devices further comprises creating corresponding device nodes in the device tree (918).

15

35. The method of claim 34, wherein said sub-node initialization operations further comprise defining resources requirements for each child driver serviceable device (918).

20

36. The method of claim 35, wherein said sub-node initialization operations further comprise:

a12. binding a child driver with a corresponding device in said device tree (920).

25

37. The method of claim 36, wherein said sub-node initialization special operations further comprise:

a13. initializing the child driver (930), which results into a nested implementation (940) of said sub-node initialization operations when said child driver is itself an intermediate driver.

30

38. The method of claim 27, wherein, for at least other drivers than the intermediary ones, said driver special operation of initializing a driver instantiation for the bound device further comprises registering the new driver instantiation in the device registry (998).

35

39. The method of claim 31, further comprising the step of:

f. providing the operating system of the computer with code (207) for launching said sub-node operations in a driver bound to a given device in the device tree.

5 40. The method of claim 39, wherein said code (207) for launching said sub-node operations of step f. comprises a synchronizing thread in the nucleus of said operating system.

10 41. The method of claim 39, wherein step f. is implemented upon computer start with a driver bound to the main bus.

42. The method of claim 39, wherein step f. is implemented during computer operation, upon detection of a new driver, with a driver bound to the main bus.

15 43. The method of claim 39, wherein step f. is implemented during computer operation, upon detection of a new device by an intermediate device, with the driver corresponding to the intermediate device.

20 44. The method of claim 39, further comprising the step of:  
g. providing the operating system of the computer with code (207) operable upon detection of a device removal to launch de-activation of the corresponding driver instantiation.

25 45. A driver, having a global driver section (6) comprising driver program code (60) and global driver data (61), both defining driver operations associated with a corresponding class of devices, and a driver initialization operation 30 enabling the construction of driver instantiations comprising local driver data, having an internal pointer to the instantiated driver.

46. A driver according to claim 45, wherein said local driver data further comprises a driver registry entry (711).

47. A driver according to claim 45, wherein said global driver section (6) further defines an operation of self-registration in a driver register.

48. A driver according to claim 45, wherein said global driver section (6) further defines an operation of probing driver serviceable devices.

5 49. An operating system, having device and driver management code, wherein said device and driver management code (207) comprises a dedicated synchronizing thread, code for tree management, code for managing a driver registry, code for managing a device registry, and code for launching selected 10 operations in a designated driver.

50. The operating system of claim 49, wherein the designated driver is the main bus driver.

1 / 13

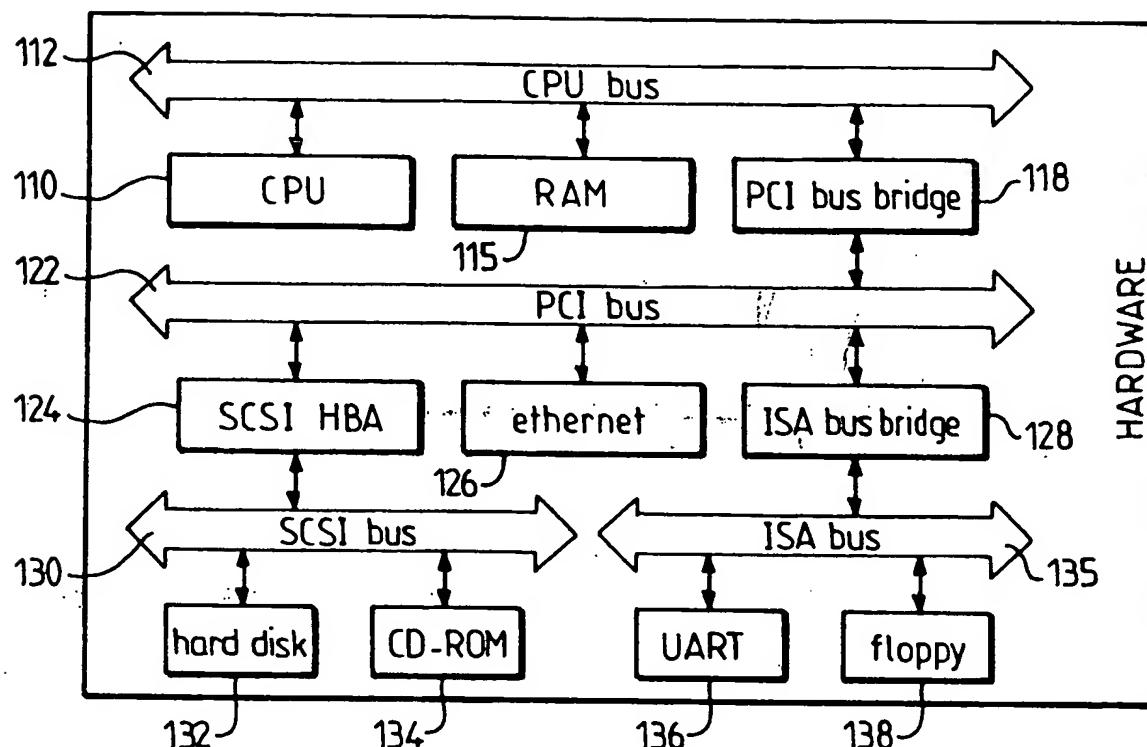


FIG.1

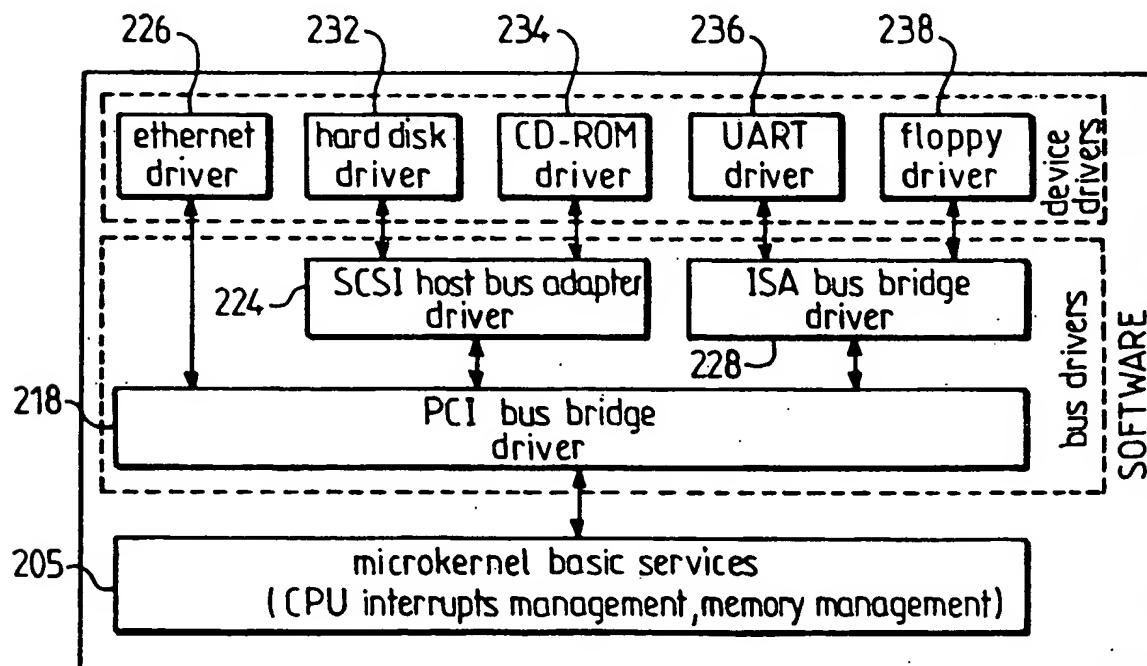


FIG.2

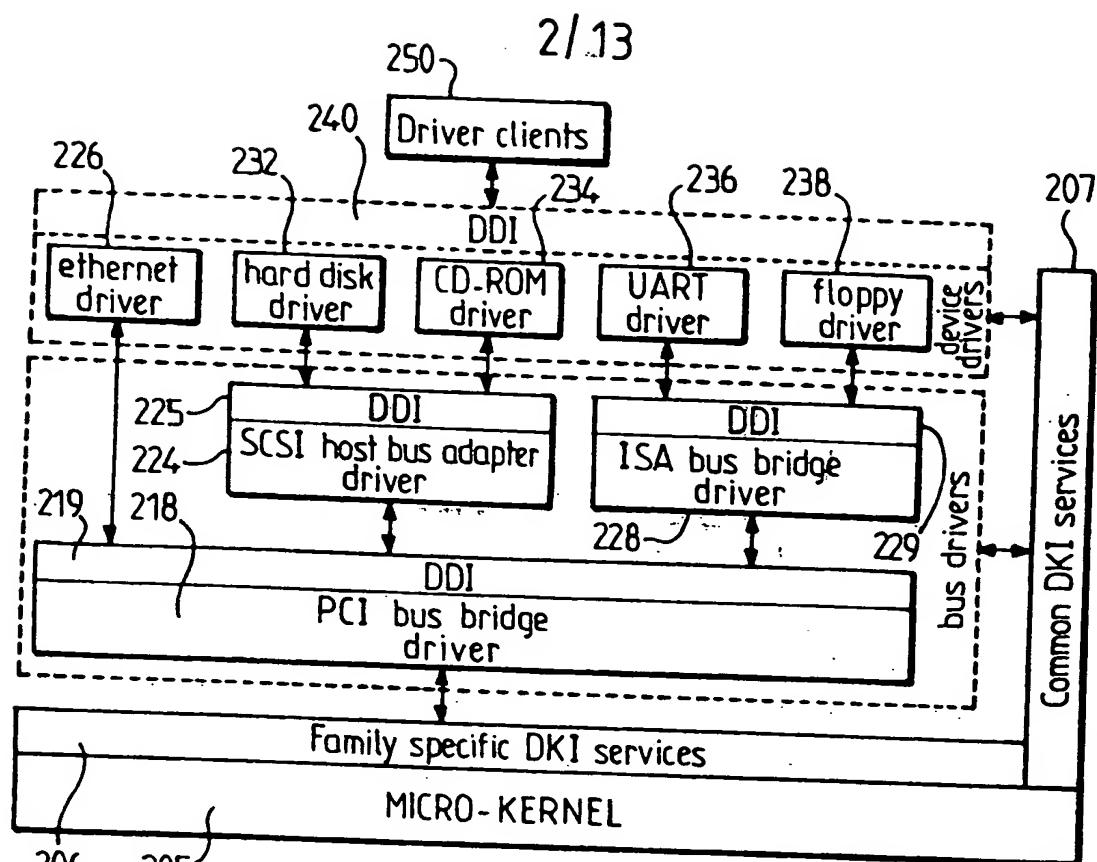


FIG. 3

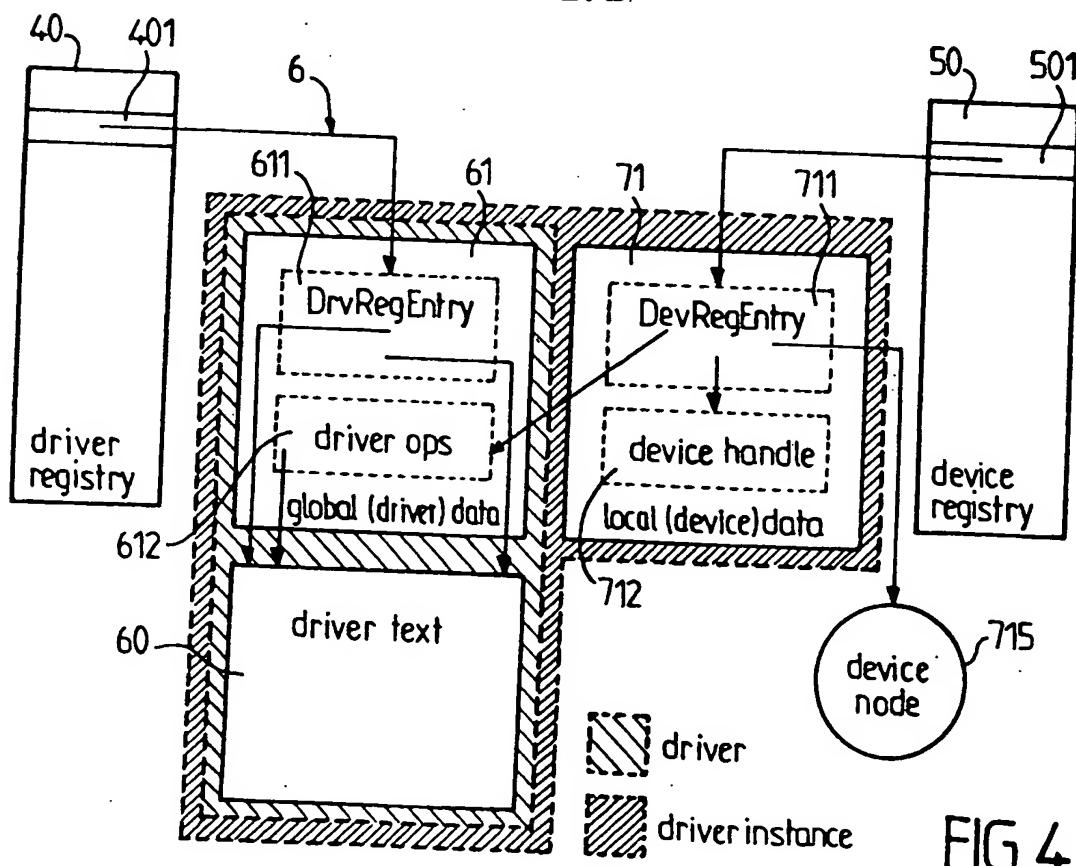


FIG. 4

3/13

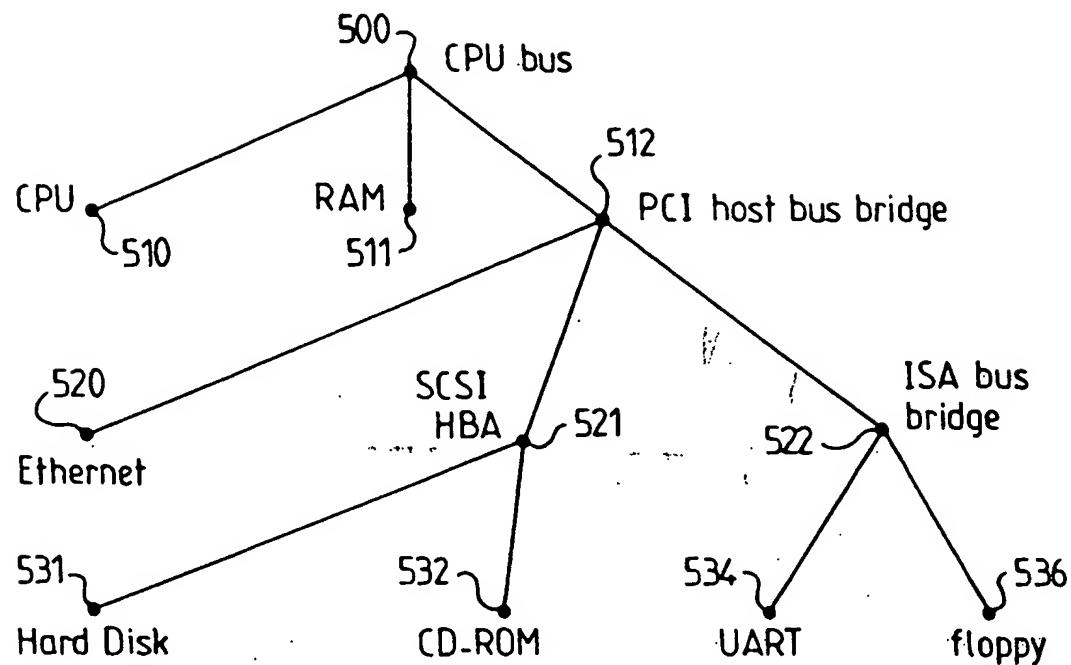


FIG.5

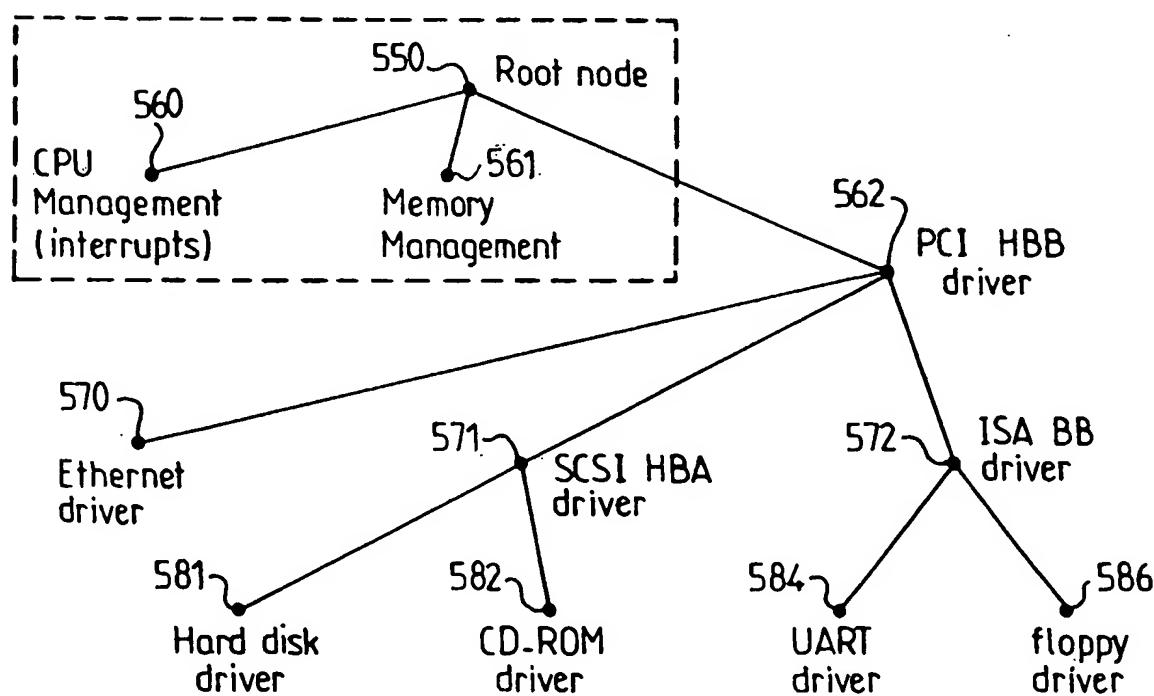


FIG.6

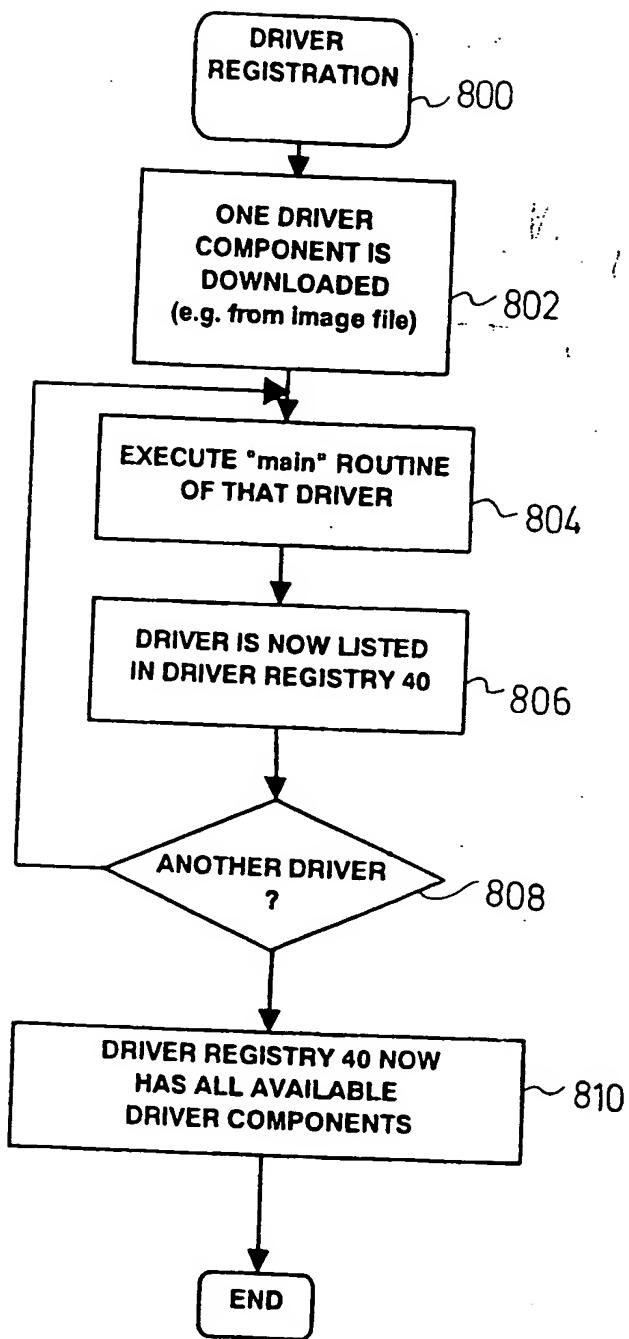


FIG. 7

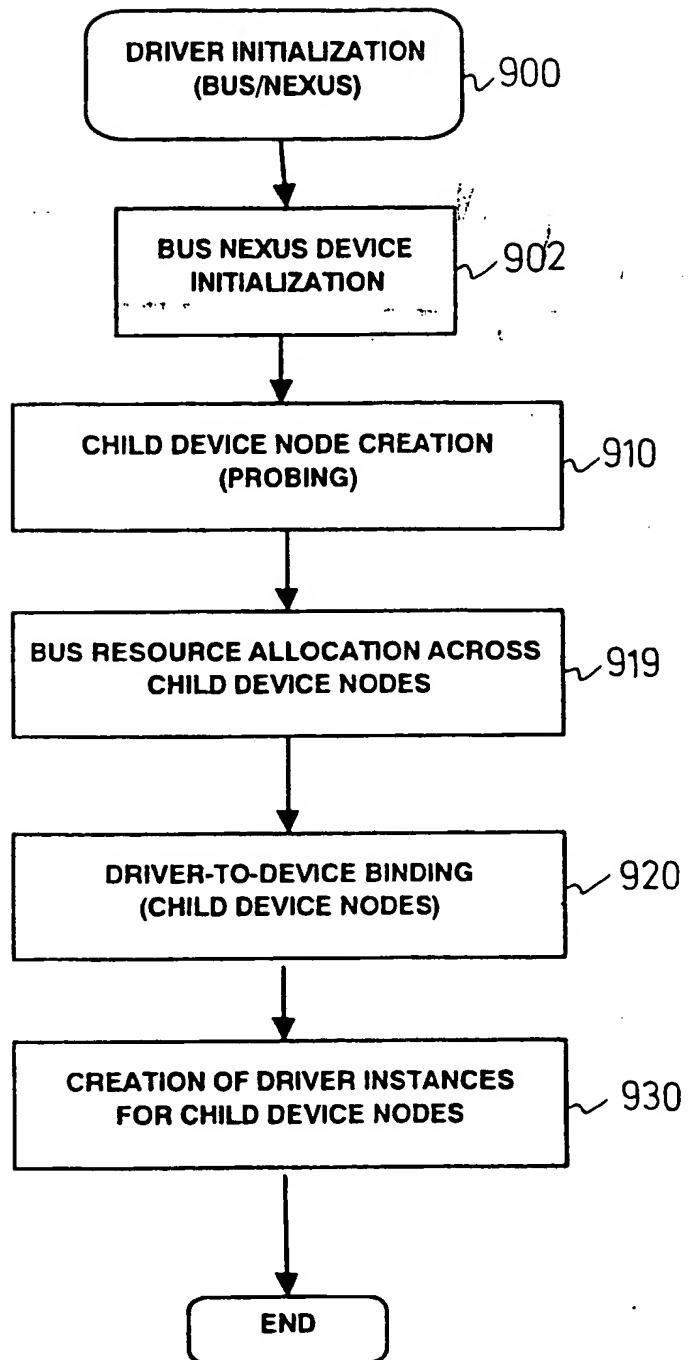


FIG. 8

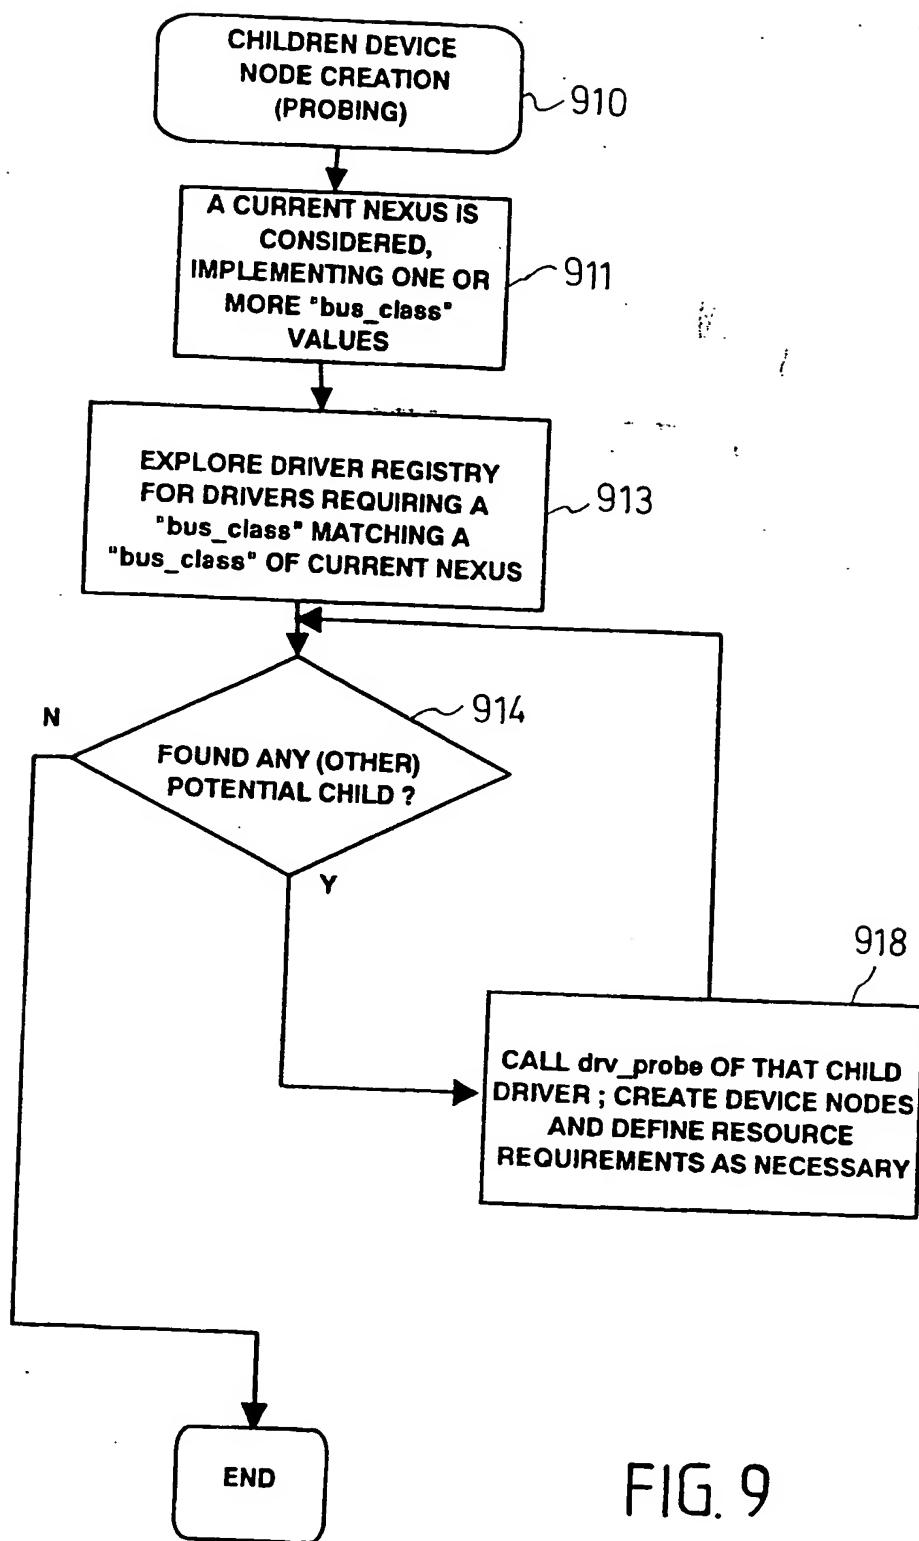


FIG. 9

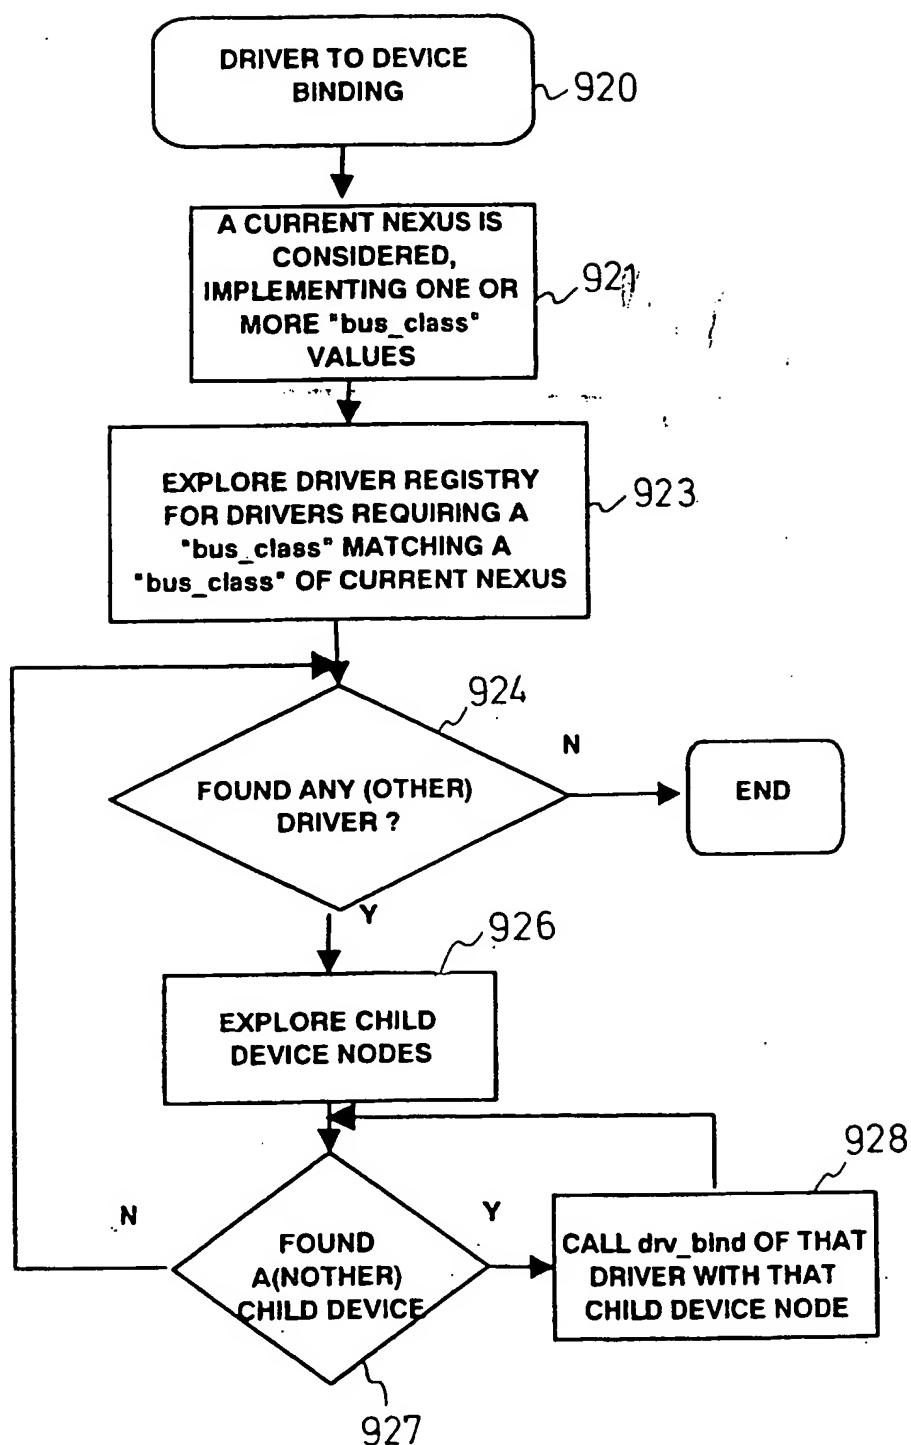


FIG. 10

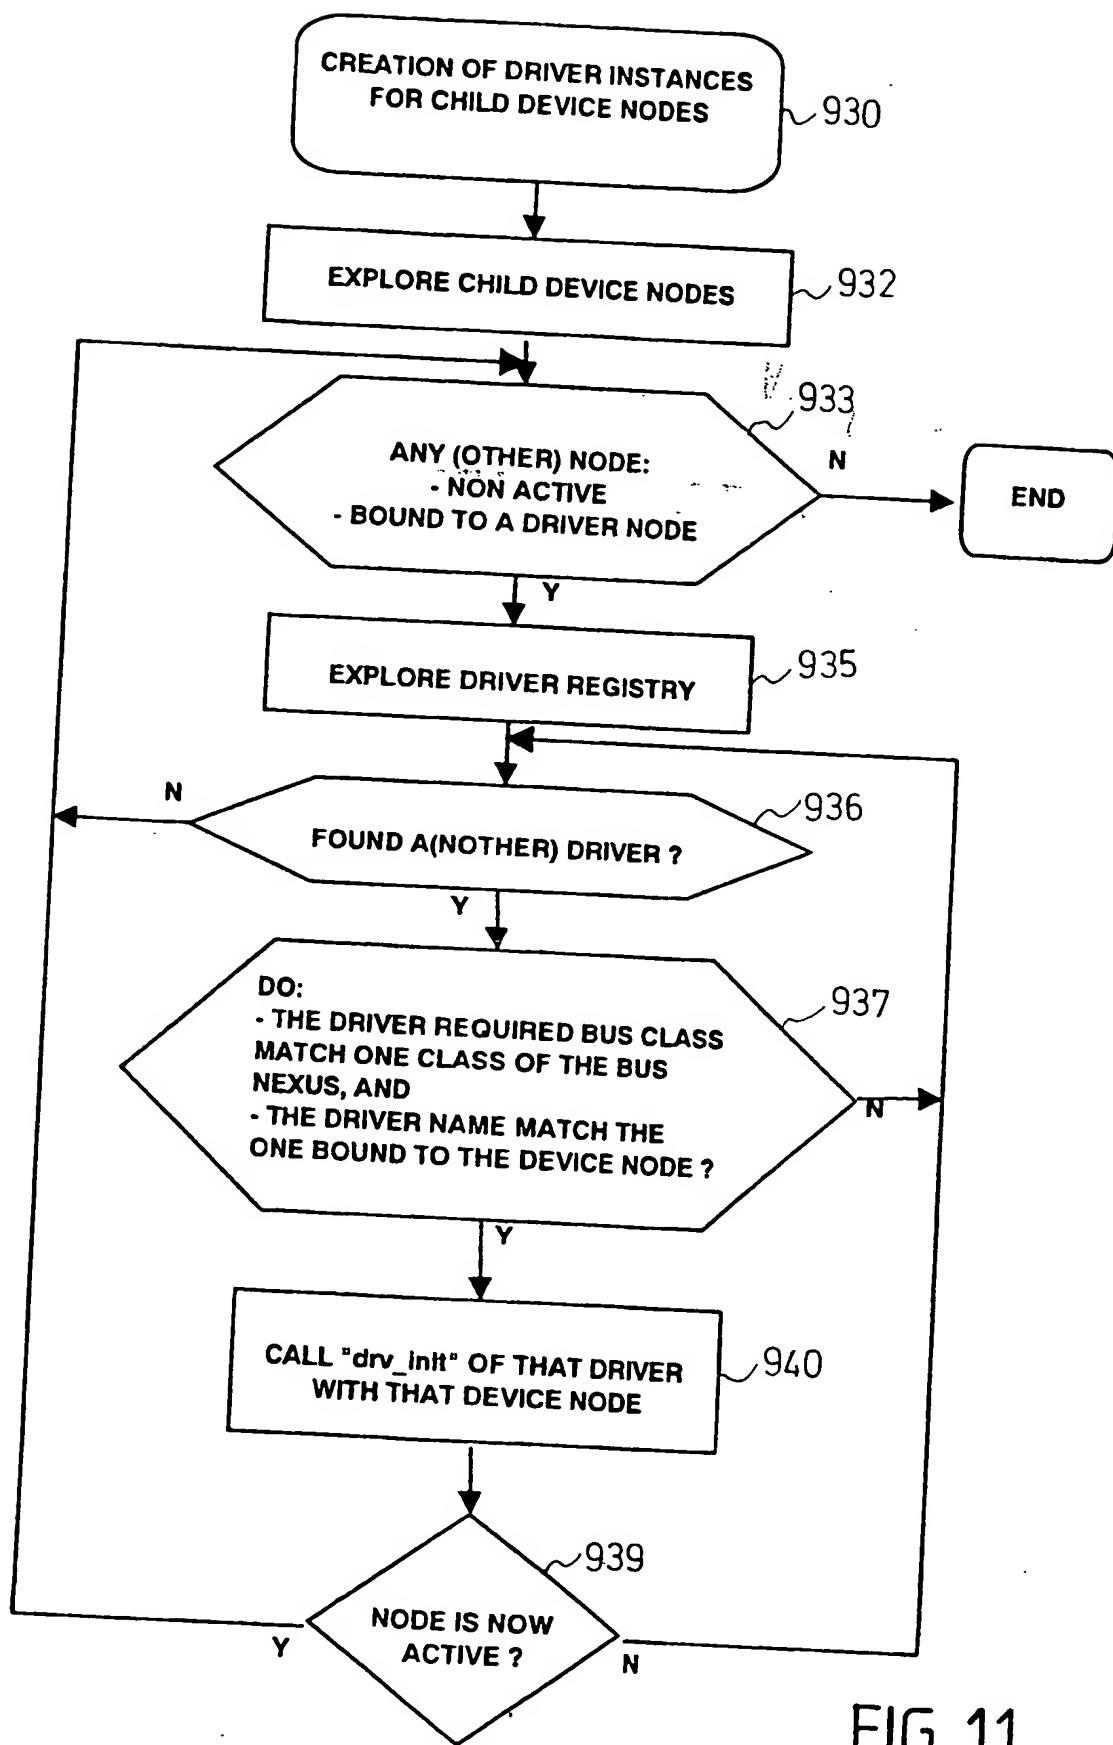


FIG. 11

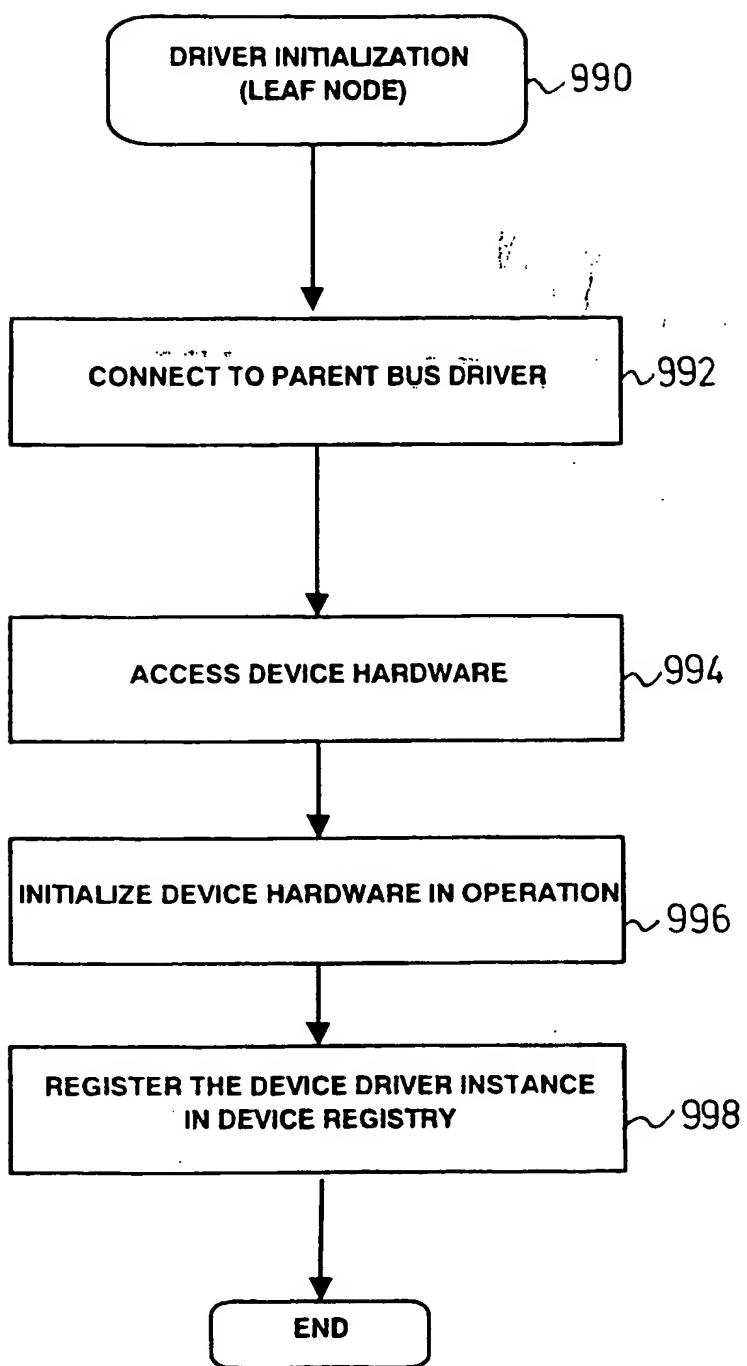


FIG. 12

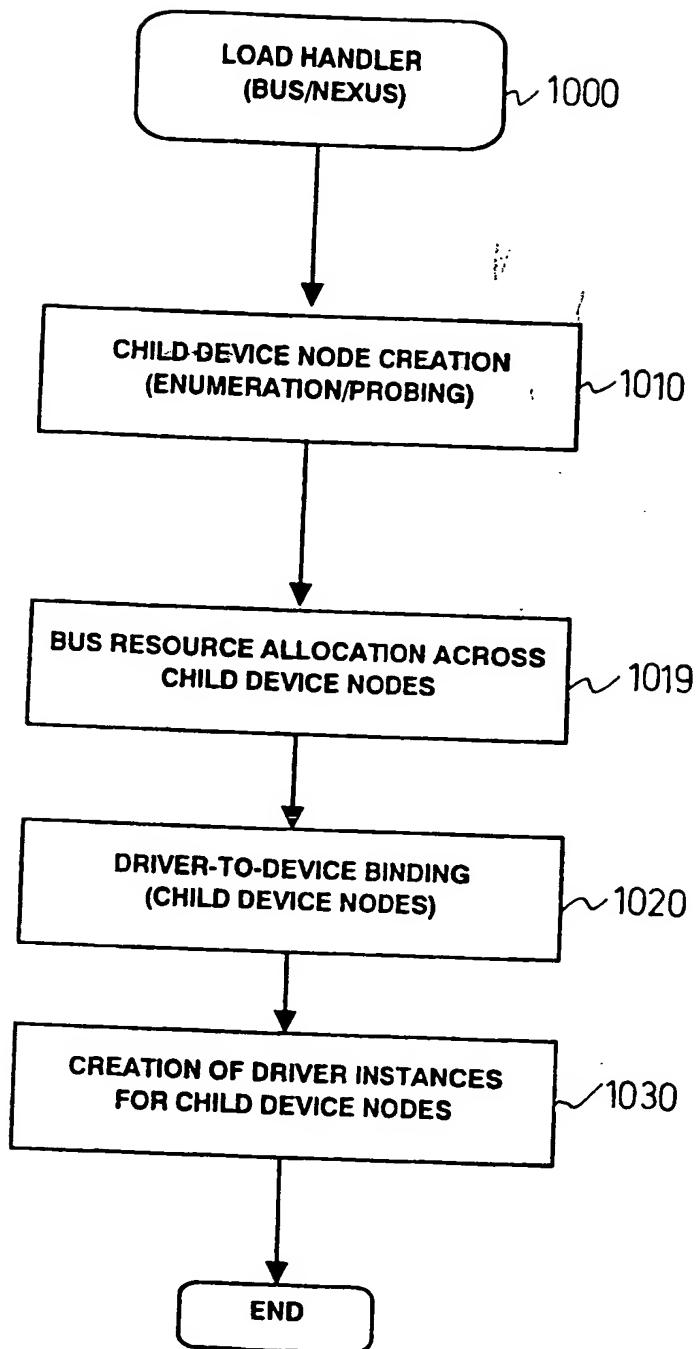


FIG. 13

Services	Base level	DKI thread	Interrupt	Blocking
dTreeNodeRoot	-	+	-	-
dTreeNodeChild	-	+	-	-
dTreeNodePeer	-	+	-	-
dTreeNodeParent	-	+	-	-
dTreeNodeAlloc	+	+	-	+
dTreeNodeFree	+	+	-	+
dTreeNodeAttach	-	+	-	-
dTreeNodeDetach	-	+	-	-
dTreePropFind	-	+	-	-
dTreePropFindNext	-	+	-	-
dTreePropLength	-	+	-	-
dTreePropValue	-	+	-	-
dTreePropName	-	+	-	-
dTreePropAlloc	+	+	-	+
dTreePropFree	+	+	-	+
dTreePropAttach	-	+	-	-
dTreePropDetach	-	+	-	-
dTreeNodeAdd	-	+	-	+
dTreeNodeFind	-	+	-	-
dTreePropAdd	-	+	-	+

FIG. 15 - TABLE T2

Services	Base level	DKI thread	Interrupt	Blocking
svDkiThreadCall	+	+	-	+
svDkiThreadTrigger	+	+	+	-

FIG. 14 - TABLE T1

Services	Base level	DKI thread	Interrupt	Blocking
svDriverRegister	+	+	-	+
svDriverLookupFirst	+	+	-	+
svDriverLookupNext	+	+	-	+
svDriverRelease	+	+	-	+
svDriverEntry	+	+	-	-
svDriverCap	+	+	-	-
svDriverUnregister	+	+	-	+

FIG. 16 - TABLE T3

Services	Base level	DKI thread	Interrupt	Blocking
svDeviceAlloc	+	+	-	+
svDeviceFree	+	+	-	+
svDeviceRegister	+	+	-	+
svDeviceUnregister	+	+	-	+
svDeviceEvent	+	+	+	-
svDeviceLookup	+	+	-	+
svDeviceRelease	+	+	-	+
svDeviceEntry	+	+	-	-

FIG.17 - TABLE T4

actions	SYS SHUTDOWN	DEV SHUTDOWN	DEV REMOVAL
notify child drivers	+	+	+
notify driver clients	-	+	+
abort operations in progress	-	-	+
reset hardware	+	-	-

FIG.18 - TABLE T5

actions	DEV SHUTDOWN	DEV REMOVAL
reset hardware	+	-
release system resources	+	+
close connection to the parent driver	+	+

FIG.19 - TABLE T6

## INTERNATIONAL SEARCH REPORT

International Application No

PCT/IB 99/02006

## A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F9/445

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, INSPEC, IBM-TDB, WPI Data

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 802 365 A (BANKS JANO ET AL) 1 September 1998 (1998-09-01)	1-7, 17, 19, 23-29, 45, 46 8, 9, 12, 16, 18, 20-22, 30, 31, 34, 38, 47-49
A	abstract column 2, line 36 - line 65 column 6, line 8 - line 44 column 7, line 45 -column 9, line 14 column 10, line 56 -column 15, line 60 column 26, line 29 - line 64 column 31, line 20 - line 49 --- -/-	

 Further documents are listed in the continuation of box C. Patent family members are listed in annex.

## \* Special categories of cited documents :

- \*'A' document defining the general state of the art which is not considered to be of particular relevance
- \*'E' earlier document but published on or after the international filing date
- \*'L' document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- \*'O' document referring to an oral disclosure, use, exhibition or other means
- \*'P' document published prior to the international filing date but later than the priority date claimed

- \*'T' later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- \*'X' document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- \*'Y' document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- \*'8' document member of the same patent family

Date of the actual completion of the international search

15 January 2001

Date of mailing of the international search report

19/01/2001

## Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

## Authorized officer

Archontopoulos, E

## INTERNATIONAL SEARCH REPORT

In  
PCT/IB 99/02006

## C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 819 107 A (ENSTROM MARK R ET AL) 6 October 1998 (1998-10-06)  abstract column 4, line 6 - line 52 column 5, line 23 - line 28 column 8, line 54 -column 9, line 45 column 25, line 15 - line 19 column 27, line 24 -column 29, line 35 figures 7A,8  ---  US 5 809 331 A (STAATS ERIK P ET AL) 15 September 1998 (1998-09-15)  abstract column 1, line 64 -column 2, line 6 column 4, line 35 -column 6, line 67  ---  US 5 949 979 A (HOLLOWELL II J RHOADS ET AL) 7 September 1999 (1999-09-07) abstract column 2, line 20 - line 62  -----	1-4, 23-26, 45,49  1-4, 23-26, 45,49  1,23,45, 49

## INTERNATIONAL SEARCH REPORT

Information on patent family members

In: International Application No

PCT/IB 99/02006

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
US 5802365	A 01-09-1998	US 5630076 A		13-05-1997
		AU 5725696 A		21-11-1996
		EP 0774137 A		21-05-1997
		WO 9635170 A		07-11-1996
US 5819107	A 06-10-1998	US 5787246 A		28-07-1998
		US 5655148 A		05-08-1997
		US 5748980 A		05-05-1998
		US 5793979 A		11-08-1998
		US 5809329 A		15-09-1998
		US 6003097 A		14-12-1999
US 5809331	A 15-09-1998	NONE		
US 5949979	A 07-09-1999	NONE		